

This Page Is Inserted by IFW Operations  
and is not a part of the Official Record

## **BEST AVAILABLE IMAGES**

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

**IMAGES ARE BEST AVAILABLE COPY.**

**As rescanning documents *will not* correct images,  
please do not report the images to the  
Image Problem Mailbox.**

	Type	Hits	Search Text	DBs
1	BRS	393	712/215.ccls.	USPAT
2	BRS	544	712/228,229.ccls.	USPAT
3	BRS	178	712/43.ccls.	USPAT
4	BRS	1	paraday.in.	USPAT
5	BRS	8	paraday	USPAT
6	BRS	2721	SMT	USPAT
7	BRS	76	simultaneous adj multithread\$3	USPAT

# NPL SEARCH 07/06/04

Set	Items	Description
S1	0	DEC()ALPHA()21464
S2	19	21464
File	2:INSPEC	1969-2004/Jun W4 (c) 2004 Institution of Electrical Engineers
File	6:NTIS	1964-2004/Jun W4 (c) 2004 NTIS, Intl Cpyrght All Rights Res
File	8:EI Compendex(R)	1970-2004/Jun W4 (c) 2004 Elsevier Eng. Info. Inc.
File	34:SciSearch(R)	Cited Ref Sci 1990-2004/Jun W4 (c) 2004 Inst for Sci Info
File	35:Dissertation Abs Online	1861-2004/May (c) 2004 ProQuest Info&Learning
File	65:Inside Conferences	1993-2004/Jul W1 (c) 2004 BLDSC all rts. reserv.
File	92:IHS Intl.Stds.& Specs.	1999/Nov (c) 1999 Information Handling Services
File	94:JICST-EPlus	1985-2004/Jun W2 (c)2004 Japan Science and Tech Corp(JST)
File	95:TEME-Technology & Management	1989-2004/Jun W1 (c) 2004 FIZ TECHNIK
File	99:Wilson Appl. Sci & Tech Abs	1983-2004/Jun (c) 2004 The HW Wilson Co.
File	103:Energy SciTec	1974-2004/Jun B2 (c) 2004 Contains copyrighted material
File	144:Pascal	1973-2004/Jun W4 (c) 2004 INIST/CNRS
File	202:Info. Sci. & Tech. Abs.	1966-2004/May 14 (c) 2004 EBSCO Publishing
File	233:Internet & Personal Comp. Abs.	1981-2003/Sep (c) 2003 EBSCO Pub.
File	239:Mathsci	1940-2004/Aug (c) 2004 American Mathematical Society
File	275:Gale Group Computer DB(TM)	1983-2004/Jul 05 (c) 2004 The Gale Group
File	434:SciSearch(R)	Cited Ref Sci 1974-1989/Dec (c) 1998 Inst for Sci Info
File	647:CMP Computer Fulltext	1988-2004/Jun W4 (c) 2004 CMP Media, LLC
File	674:Computer News Fulltext	1989-2004/Jun W2 (c) 2004 IDG Communications
File	696:DIALOG Telecom. Newsletters	1995-2004/Jul 05 (c) 2004 The Dialog Corp.

2/5/1 (Item 1 from file: 2)  
DIALOG(R)File 2:INSPEC  
(c) 2004 Institution of Electrical Engineers. All rts. reserv.

6561611 INSPEC Abstract Number: C2000-05-5130-036

**Title: EV8 'Arana' highlights alpha architectural advances**

Author(s): Shannon, T.

Journal: Digital Systems Report vol.22, no.1 p.22-3

Publisher: Computer Economics,

Publication Date: Spring 2000 Country of Publication: USA

CODEN: DSREFF ISSN: 1086-9638

SICI: 1086-9638(200021)22:1L:22:AHAA;1-C

Material Identity Number: F163-2000-002

Language: English Document Type: Journal Paper (JP)

Treatment: Practical (P); Product Review (R)

Abstract: Compaq has developed an alpha **21464** EV8 "Arana" CPU. The article provides a look at what alpha customers can expect to find inside their systems early in the new millennium. The author highlights the approach Compaq will take to make an EV8 uniprocessor act like a four-way SMP system. Called simultaneous multithreading (SMT), the technique exploits thread-level parallelism to make better use of processor resources without resorting to chip multiprocessing or instruction-level parallelism a la Intel's Itanium CPU. (0 Refs)

Subfile: C

Descriptors: Compaq computers; equipment evaluation; microprocessor chips ; multi-threading; processor scheduling

Identifiers: EV8 Arana; alpha architectural advances; Compaq; alpha **21464** EV8 Arana CPU; EV8 uniprocessor; four-way SMP system; simultaneous multithreading; thread-level parallelism; processor resources; Itanium CPU

Class Codes: C5130 (Microprocessor chips); C5440 (Multiprocessing systems); C6150N (Distributed systems software); C6110P (Parallel programming)

Copyright 2000, IEE

2/5/8 (Item 2 from file: 275)  
DIALOG(R)File 275:Gale Group Computer DB(TM)  
(c) 2004 The Gale Group. All rts. reserv.

02466711 SUPPLIER NUMBER: 68024916 (USE FORMAT 7 OR 9 FOR FULL TEXT)

**NETWORKING GETS XSTREAM : Startup Debuts Simultaneous Multithreading in**

**Network Processor.(Company Business and Marketing)**

Glaskowsky, Peter N.

Microprocessor Report, 14, 11, 17

Nov, 2000

ISSN: 0899-9341 LANGUAGE: English RECORD TYPE: Fulltext

WORD COUNT: 1157 LINE COUNT: 00097

COMPANY NAMES: Xstream Logic Inc.--Product development

GEOGRAPHIC CODES/NAMES: 1USA United States

DESCRIPTORS: Company product planning; Company technology development; Microprocessor

EVENT CODES/NAMES: 331 Product development

PRODUCT/INDUSTRY NAMES: 3674000 (Semiconductor Devices); 3674124 (Microprocessor Chips)

SIC CODES: 3674 Semiconductors and related devices

NAICS CODES: 334413 Semiconductor and Related Device Manufacturing

FILE SEGMENT: CD File 275

?t s2/5,k/9-19

2/5,K/9 (Item 3 from file: 275)  
DIALOG(R)File 275:Gale Group Computer DB(TM)  
(c) 2004 The Gale Group. All rts. reserv.

02409123 SUPPLIER NUMBER: 62710910 (USE FORMAT 7 OR 9 FOR FULL TEXT)

**IBM PAVING THE WAY TO 0.10 MICRON : First Copper, Then SOI, Now Low-k and**

**E-Beams.(Company Business and Marketing)**

Diefendorff, Keith

Microprocessor Report, 14, 5, 1

May, 2000

ISSN: 0899-9341      LANGUAGE: English      RECORD TYPE: Fulltext

WORD COUNT: 6553      LINE COUNT: 00532

COMPANY NAMES: International Business Machines Corp.--Research  
GEOGRAPHIC CODES/NAMES: 1USA United States  
DESCRIPTORS: Research and development; Company technology development  
EVENT CODES/NAMES: 310 Science & research  
PRODUCT/INDUSTRY NAMES: 3674000 (Semiconductor Devices)  
SIC CODES: 3674 Semiconductors and related devices  
NAICS CODES: 334413 Semiconductor and Related Device Manufacturing  
TICKER SYMBOLS: IBM  
FILE SEGMENT: CD File 275

... copper CMOS-8S, and eventually to SOI.

Also, at Microprocessor Forum last October, Compaq said that its **21464** would be constructed in a 0.13-micron copper low-k SOI process (see MPR 11/15/99-msb, "Alpha **21464** targets 1.7GHz in 2003"). Furthermore, rumors persist that Compaq is on the verge of announcing a...

...is also negotiating for access to SOI-based CMOS-8S2 and CMOS-9S for its 21364 and **21464**. Such a deal would be a good move for Compaq and would give us a more favorable...

**2/5,K/10 (Item 4 from file: 275)**

DIALOG(R)File 275:Gale Group Computer DB(TM)

(c) 2004 The Gale Group. All rts. reserv.

02377059      SUPPLIER NUMBER: 59617880      (USE FORMAT 7 OR 9 FOR FULL TEXT)

**Microprocessor Report Analysts' Choice Awards. (Buyers Guide)**

Microprocessor Report, 14, 1, NA

Jan, 2000

DOCUMENT TYPE: Buyers Guide      ISSN: 0899-9341      LANGUAGE: English

RECORD TYPE: Fulltext

WORD COUNT: 225      LINE COUNT: 00021

GEOGRAPHIC CODES/NAMES: 1USA United States  
DESCRIPTORS: Hardware buyers' guide; Microprocessor  
EVENT CODES/NAMES: 330 Product information;350 Product standards, safety, & recalls  
PRODUCT/INDUSTRY NAMES: 3674124 (Microprocessor Chips)  
NAICS CODES: 334413 Semiconductor and Related Device Manufacturing  
FILE SEGMENT: CD File 275

... 05)

Best New Technology: \* Winner: IBM POWER4 (see MPR 10/6/99-02) \*  
Honorable Mention: Compaq Alpha **21464** (see MPR 12/6/99-01) \* HAL SPARC64 V (see MPR 11/15/99-01) \* HP/Intel...

**2/5,K/11 (Item 5 from file: 275)**

DIALOG(R)File 275:Gale Group Computer DB(TM)

(c) 2004 The Gale Group. All rts. reserv.

02377058      SUPPLIER NUMBER: 59617879      (USE FORMAT 7 OR 9 FOR FULL TEXT)

**BEST NEW TECHNOLOGY: POWER4 : IBM's Chip Multiprocessor Is Analysts' Choice for Technology Award. (Product Information)**

Diefendorff, Keith

Microprocessor Report, 14, 1, NA

Jan, 2000

ISSN: 0899-9341      LANGUAGE: English      RECORD TYPE: Fulltext

WORD COUNT: 2694      LINE COUNT: 00210

COMPANY NAMES: Compaq Computer Corp.--Achievements and awards;  
International Business Machines Corp.--Products  
GEOGRAPHIC CODES/NAMES: 1USA United States

DESCRIPTORS: Product description/specification; Microprocessor  
EVENT CODES/NAMES: 330 Product information  
PRODUCT/INDUSTRY NAMES: 3674124 (Microprocessor Chips)  
NAICS CODES: 334413 Semiconductor and Related Device Manufacturing  
TRADE NAMES: Compaq Alpha 21464 (Microprocessor)--Achievements and  
awards; IBM Power4 (Microprocessor)--Achievements and awards  
FILE SEGMENT: CD File 275

TEXT:

...annual Microprocessor Report Technology Award was given to IBM,  
beating out five other nominees: Compaq's Alpha 21464, HAL's SPARC64 V,  
HP and Intel's IA-64 architecture, Sony and Toshiba's Emotion Engine...  
... for first place was so tight we also awarded an Honorable Mention  
to Compaq for the Alpha 21464.

The "Unlimited-Class" Award

At our dinner meeting on January 27, MDR analysts presented four  
Analysts' Choice...

...future.

Our first nomination for this year's Microprocessor Report Technology  
Award went to the Compaq Alpha 21464 -for its adoption of simultaneous  
multithreading, or SMT. This clever idea originated with Susan Eggers and  
Hank...

...of which is unknown at this time. It was primarily this uncertainty that  
knocked SMT and the 21464 out of first place for our Technology Award.

Unbelievable Horsepower in a Kid's Game

For packing...

TRADE NAMES: Compaq Alpha 21464 (Microprocessor...

2/5,K/12 (Item 6 from file: 275)

DIALOG(R) File 275:Gale Group Computer DB(TM)

(c) 2004 The Gale Group. All rts. reserv.

02366363 SUPPLIER NUMBER: 58926964 (USE FORMAT 7 OR 9 FOR FULL TEXT)

**The birth of a NEW processor. (Industry Trend or Event)**

Britt, Russ

Electronic Business, 26, 1, 62

Jan, 2000

ISSN: 0163-6197

LANGUAGE: English

RECORD TYPE: Fulltext; Abstract

WORD COUNT: 3216 LINE COUNT: 00263

ABSTRACT: John Huck of Hewlett-Packard and John Crawford of Intel joined  
forces to create the microprocessor dubbed Itanium. The chip, which is not  
yet out is thought to be underpowered by some in the industry. there is  
also concern about how much support the comp;companies who developed this  
chip will give each other. This article covers industry concerns and the  
background of the makings of Itanium.

COMPANY NAMES: Hewlett-Packard Co.--Management; Intel Corp.--Management

GEOGRAPHIC CODES/NAMES: 1USA United States

DESCRIPTORS: Company business management; Microprocessor; Industry trend

EVENT CODES/NAMES: 220 Strategy & planning

PRODUCT/INDUSTRY NAMES: 3573000 (Computers & Peripherals); 3674000

(Semiconductor Devices); 3674124 (Microprocessor Chips)

NAICS CODES: 334111 Electronic Computer Manufacturing; 334413

Semiconductor and Related Device Manufacturing

TRADE NAMES: Intel Itanium (Microprocessor)--Planning

FILE SEGMENT: CD File 275

... Alpha 21364 Late 2000 New system interface; 1.6-  
GHz version due in early 2002  
Compaq Alpha 21464 2002 New multithreaded core  
Sun UltraSparc-3 First half 2001 600-MHz at 0.18 micron  
Sun...

2/5,K/13 (Item 7 from file: 275)  
DIALOG(R)File 275:Gale Group Computer DB(TM)  
(c) 2004 The Gale Group. All rts. reserv.

02362763 SUPPLIER NUMBER: 58500786 (USE FORMAT 7 OR 9 FOR FULL TEXT)  
**X86 Outdoes RISC Performance. (Company Business and Marketing)**  
Gwennap, Linley  
Microprocessor Report, 13, 17, NA  
Dec 27, 1999  
ISSN: 0899-9341 LANGUAGE: English RECORD TYPE: Fulltext  
WORD COUNT: 3865 LINE COUNT: 00294

COMPANY NAMES: Intel Corp.--Product development; Advanced Micro Devices  
Inc.--Product development; MIPS Computer Systems Inc.--Product  
development  
GEOGRAPHIC CODES/NAMES: 1USA United States  
DESCRIPTORS: Review of past year; Preview of coming year; Company  
technology development  
EVENT CODES/NAMES: 331 Product development  
PRODUCT/INDUSTRY NAMES: 3674124 (Microprocessor Chips)  
NAICS CODES: 334413 Semiconductor and Related Device Manufacturing  
FILE SEGMENT: CD File 275

... GHz 21264. In addition, Compaq earns the Biggest Crystal Balls  
award for forecasting a 1.7-GHz **21464** in early 2003.

The 21264 remains the yardstick for measuring other high-end  
processors, leading the pack...

2/5,K/14 (Item 8 from file: 275)  
DIALOG(R)File 275:Gale Group Computer DB(TM)  
(c) 2004 The Gale Group. All rts. reserv.

02358406 SUPPLIER NUMBER: 58248504 (USE FORMAT 7 OR 9 FOR FULL TEXT)  
**Compaq Chooses SMT for Alpha : Simultaneous Multithreading Exploits  
Instruction- and Thread-Level Parallelism. (Compaq Alpha 21464 ) (Product  
Information)**  
Diefendorff, Keith  
Microprocessor Report, 13, 16, NA  
Dec 6, 1999  
ISSN: 0899-9341 LANGUAGE: English RECORD TYPE: Fulltext  
WORD COUNT: 5376 LINE COUNT: 00433

COMPANY NAMES: Compaq Computer Corp.--Products  
GEOGRAPHIC CODES/NAMES: 1USA United States  
DESCRIPTORS: Product description/specification; Microprocessor; Processor  
architecture  
EVENT CODES/NAMES: 330 Product information  
PRODUCT/INDUSTRY NAMES: 3674124 (Microprocessor Chips)  
NAICS CODES: 334413 Semiconductor and Related Device Manufacturing  
TRADE NAMES: Compaq Alpha **21464** (Microprocessor)--Design and  
construction  
FILE SEGMENT: CD File 275  
**Compaq Chooses SMT for Alpha : Simultaneous Multithreading Exploits  
Instruction- and Thread-Level Parallelism. (Compaq Alpha 21464 ) (Product  
Information)**

TEXT:

...within Compaq. His efforts have apparently paid off, as Compaq has  
officially adopted SMT for the Alpha **21464** (see MPR 11/15/99, p. 13),  
code-named EV8, which is due to appear in systems...

TRADE NAMES: Compaq Alpha **21464** (Microprocessor...

2/5,K/15 (Item 9 from file: 275)  
DIALOG(R)File 275:Gale Group Computer DB(TM)

(c) 2004 The Gale Group. All rts. reserv.

02349930 SUPPLIER NUMBER: 57588372 (USE FORMAT 7 OR 9 FOR FULL TEXT)  
**Alpha 21464 Targets 1.7 GHz in 2003. (Compaq details plans for Alpha EV8 processor) (Company Business and Marketing)**  
Microprocessor Report, 13, 15, NA  
Nov 15, 1999  
ISSN: 0899-9341 LANGUAGE: English RECORD TYPE: Fulltext  
WORD COUNT: 489 LINE COUNT: 00040

COMPANY NAMES: Compaq Computer Corp.--Product development  
GEOGRAPHIC CODES/NAMES: 1USA United States  
DESCRIPTORS: Microprocessor; Company product planning  
EVENT CODES/NAMES: 331 Product development  
PRODUCT/INDUSTRY NAMES: 3674124 (Microprocessor Chips)  
NAICS CODES: 334413 Semiconductor and Related Device Manufacturing  
TRADE NAMES: Compaq Alpha EV-8 (Microprocessor)--Product development  
FILE SEGMENT: CD File 275

**Alpha 21464 Targets 1.7 GHz in 2003. (Compaq details plans for Alpha EV8 processor) (Company Business and Marketing)**

TEXT:

Determined to maintain leadership performance well into the next century, Compaq disclosed plans for its futuristic **21464** processor at last month's Microprocessor Forum.

... is not scheduled to appear in systems until early 2003.

According to Compaq's Joel Emer, the **21464** will achieve single-thread performance leadership using an eight-way superscalar processor core running at speeds of...

...as copper, low-k dielectrics, and SOI. Compaq did not name the fab, but we expect the **21464**, like current Alpha chips, will be built by Samsung and possibly another foundry. Emer said the design...

...and up, but these speeds will require more advanced IC process technology.

To further boost performance, the **21464** will implement four virtual processors on the chip, using a technique called simultaneous multithreading (SMT). This method...

...much as ?in a multi-threaded environment, as is common in servers.

The system interface of the **21464** will be similar to that of the 21364 (see MPR 10/26/98, p. 12), which is due to appear in systems in early 2001. Like that chip, the **21464** will have a large on-chip L2 cache, several Rambus channels for main memory, and four additional...

...to accommodate the more powerful core, but Compaq declined to provide additional details.

In 2002, the **21464** will compete against Intel's Madison, a 0.13-micron version of McKinley. We expect Madison to...

...on single-threaded programs.

Both chips are likely to deliver in excess of 130 SPECint95 (base). The **21464**, however, could have an edge in servers, due to its multithreaded design; we expect the McKinley/ Madison...

2/5,K/16 (Item 10 from file: 275)

DIALOG(R)File 275:Gale Group Computer DB(TM)  
(c) 2004 The Gale Group. All rts. reserv.

02338696 SUPPLIER NUMBER: 56025118 (USE FORMAT 7 OR 9 FOR FULL TEXT)  
**Power4 Focuses on Memory Bandwidth. (new IBM architecture) (Product Development)**  
Diefendorff, Keith  
Microprocessor Report, 13, 13, NA  
Oct 6, 1999



ISSN: 0899-9341      LANGUAGE: English      RECORD TYPE: Fulltext  
WORD COUNT: 6484      LINE COUNT: 00501

COMPANY NAMES: IBM Personal Computer Co.--Product development  
GEOGRAPHIC CODES/NAMES: 1USA United States  
DESCRIPTORS: Hardware product development; Microprocessor  
EVENT CODES/NAMES: 331 Product development  
PRODUCT/INDUSTRY NAMES: 3674124 (Microprocessor Chips)  
NAICS CODES: 334413 Semiconductor and Related Device Manufacturing  
TRADE NAMES: IBM Power4 (Microprocessor)--Product development  
FILE SEGMENT: CD File 275

... raw bandwidth numbers. Since the topologies are different, however, the bandwidth numbers are difficult to compare.

The 21464, due out sometime in 2002, will be a multithreaded version of a new core, designed to exploit...

2/5,K/17      (Item 11 from file: 275)  
DIALOG(R)File 275:Gale Group Computer DB(TM)  
(c) 2004 The Gale Group. All rts. reserv.

02126164      SUPPLIER NUMBER: 20062656      (USE FORMAT 7 OR 9 FOR FULL TEXT)  
**DEC PLOTS ALPHA RISC PROGRESS TO 2003 AND BEYOND.**  
Computergram International, n3311, pCGN12120001  
Dec 12, 1997  
ISSN: 0268-716X      LANGUAGE: English      RECORD TYPE: Fulltext  
WORD COUNT: 517      LINE COUNT: 00041

FILE SEGMENT: CD File 275

TEXT:

...cram more transistors onto the chips - although the first chip to use .18 micron - the EV8, or 21464, still won't actually be delivered until 2001. It will now have 18 million, 100 nanosecond transistors...

2/5,K/18      (Item 1 from file: 647)  
DIALOG(R)File 647:CMP Computer Fulltext  
(c) 2004 CMP Media, LLC. All rts. reserv.

01168572      CMP ACCESSION NUMBER: EET19980803S0008  
**Technical details emerge on code-optimization schemes for Merced, Alpha 21364 - Intel, Compaq gird for 64-bit MPU face-off**  
Alexander Wolfe  
ELECTRONIC ENGINEERING TIMES, 1998, n 1019, PG1  
PUBLICATION DATE: 980803  
JOURNAL CODE: EET      LANGUAGE: English  
RECORD TYPE: Fulltext  
SECTION HEADING: News  
WORD COUNT: 1667  
TEXT:

Santa Clara, Calif. - A battle is heating up at the bleeding edge of microprocessor technology as Intel Corp. and Compaq Computer Corp.'s Alpha group rush to ready their competing 64-bit architectures. New technical details have come to light about the race, which pits Intel's Merced, due out in mid-2000, against the next-generation Alpha CPU, known as the 21364. Compaq acquired the Alpha design team when it bought Digital Equipment Corp. in June.

COMPANY NAMES (DIALOG GENERATED): Compaq Computer Corp ; Digital Equipment Corp ; Digital Palo Alto Design Center ; EE Times ; Intel Corp ; IA ; Microprocessor Forum ; Speeds ; Texas Instruments Inc ; University of Illinois at Urbana Champaign

... as the 364 effort proceeds, the Alpha team on the East Coast is beginning work on the 21464. Interestingly, that device, not the 364, is the first Alpha chip slated to use 0.18-micron...

2/5,K/19 (Item 2 from file: 647)  
DIALOG(R)File 647:CMP Computer Fulltext  
(c) 2004 CMP Media, LLC. All rts. reserv.

00521605 CMP ACCESSION NUMBER: CRN19921130S0843

**Emerging-brand monitors roll out during shortages**

MICHELLE GRAZIOSE

COMPUTER RESELLER NEWS, 1992, n 502, 143

PUBLICATION DATE: 921130

JOURNAL CODE: CRN LANGUAGE: English

RECORD TYPE: Fulltext

SECTION HEADING: SOURCING

WORD COUNT: 676

TEXT:

Las Vegas

Smack in the middle of the worldwide shortage of 14-inch SVGA monitors, at least two new Taiwanese emerging-brand display vendors introduced new lines at Comdex/Fall this month.

... also display the same parameter confirmation information on the front series of vertical LED bars.

The CA- 21464 , 14-inch model features a non-glare screen; 0.28-millimeter dot pitch; maximum resolution of 1...

2/9/8 (Item 2 from file: 275)  
DIALOG(R) File 275:Gale Group Computer DB(TM)  
(c) 2004 The Gale Group. All rts. reserv.

02466711 SUPPLIER NUMBER: 68024916 (THIS IS THE FULL TEXT)  
**NETWORKING GETS XSTREAM : Startup Debuts Simultaneous Multithreading in  
Network Processor.(Company Business and Marketing)**  
Glaskowsky, Peter N.  
Microprocessor Report, 14, 11, 17  
Nov, 2000  
ISSN: 0899-9341 LANGUAGE: English RECORD TYPE: Fulltext  
WORD COUNT: 1157 LINE COUNT: 00097

TEXT:

At last month's Microprocessor Forum , XStream Logic described a new CPU architecture the company is developing for high-level network processing. The basis of the new architecture is simultaneous multithreading (SMT), a technique that first appeared at the Forum last year in Compaq's presentation on the Alpha **21464** (see MPR 12/6/99-01, "Compaq Chooses SMT for Alpha"). (XStream prefers to use the term "Dynamic Multistreaming," or DMS, instead of SMT.)

In some ways, the XStream effort is more impressive. Compaq plans to support four simultaneous threads in hardware; XStream will support eight. The Alpha chip is aimed at the high end of the server and workstation market, which can support very high chip prices; XStream is targeting more cost-sensitive networking products. And while the **21464** is not likely to appear until late in 2002, we expect XStream to announce chips next year.

XStream's DMS technology will be applied to what the company calls a "MIPS-like" instruction-set architecture. XStream believes this approach will provide a simpler and more efficient software-development environment than competing parts that use chip multiprocessing (CMP), VLIW, or wide superscalar cores. These alternatives require more-complex development tools to handle inter-processor communication and extract parallelism from code written in high-level languages. The XStream approach, on the other hand, is consistent with the single-CPU, multithreaded programming model that has been used for years in operating systems such as Unix and Windows NT.

New Core Supports Eight Threads

XStream says its first DMS core will provide hardware support for eight simultaneous threads, as Figure 1 shows. Each thread has its own instruction queue and register file. Eight function units are also available, along with two address-generator units used for write operations. The dispatch unit looks at the next four instructions from each thread and dispatches up to eight of these 32 instructions to the function units. XStream has not described the algorithm it uses to select instructions for dispatch except to say that instructions from each thread are dispatched in order.

The instruction and data caches are 64K in size and four-way set-associative. The instruction cache has a 64-byte line size, while the data cache has a 32-byte line size. In each clock period, 16 instructions from each of up to two threads can be transferred from the instruction cache to the instruction queues. Instructions in these queues may be reused to reduce the branch penalty for short branches. Also, during each clock period, not more than two write and two read operations can be completed to the data cache.

At nine stages, the DMS pipeline is deeper than those found in most MIPS cores. This depth is due in part to the extra control logic required for multithreading. As Figure 2 shows, the DMS core requires separate stages to select and queue instructions, reflecting the complexity added by the multiple instruction queues. The Dispatch stage decouples the two halves of the pipeline; instructions wait in this stage until the necessary execution resources are available and are then dispatched in program order. Register reads also take place during the Dispatch stage. The Memory stage handles data-cache reads for load instructions or register writes for register-to-register ALU and transfer instructions. Register writes for load instructions that hit in the data cache are performed in the Write stage. The Store stage is used to complete data-cache writes. Some of these stages can be skipped when not needed, reducing the effective pipeline

length.

XStream has said little about the instruction set it will support, other than the fact that it will be MIPS-like and include some networking-specific enhancements. The company has also been silent about its plans for external interfaces, such as the CPU bus and the memory controller.

#### Core Augmented by Support Functions

XStream's processors will include a full MMU. This feature is required by most of today's multithreaded operating systems. XStream is also developing a packet-management unit (PMU) to handle simple packet processing functions such as packet-memory allocation and deallocation, garbage collection, byte gathering, and network-interface I/O with only minimal processor supervision. Unlike similar peripherals on existing network processors, the PMU will also have access to the context registers within the DMS core, allowing it to set up critical pointers and data values before threads begin execution.

XStream's decision to offload these functions from the CPU core further emphasizes the company's focus on the networking market. Even without its unspecified networking instructions, the DMS core would be useful in other markets, such as consumer electronics, where its straightforward programming model and high throughput would be a good match to conventional software-development practices.

XStream will focus on application-specific functions such as load balancing and content-based filtering in the higher layers of the OSI networking model. These functions require a processor architecture that works well on highly conditional code, a requirement that rules out most of today's existing network processors that were designed for lower-level networking functions like switching and routing. Many of these high-level functions operate on streams of data, not just individual packets. The interthread communication and synchronization functions required to process these streams are easier to implement on the single-core XStream architecture than on competing multiprocessor network-processor architectures.

Although XStream has not released estimates for clock speeds, die sizes, or other implementation details, the company is setting its sights high. XStream expects to deliver processors suitable for handling Internet data at speeds up to 10Gb/s--the speed of an OC-192 fiber-optic link.

To succeed, XStream must deliver fast chips and all the other hardware and software components required to create complete network-processing systems. This is a big task, and XStream is still a small company. XStream has already announced a relationship with MontaVista Software, which is doing a version of Linux for the XStream architecture, and other such relationships are being developed.

XStream has yet to demonstrate the superiority of the DMS architecture for networking--or for any other application--but DMS is clearly unique in the crowded network-processor market. Its programming model, in particular, appears to be substantially more straightforward than those of competing products. This advantage should be enough to give XStream a chance at success, despite the intense competition it faces.

#### RELATED ARTICLE: Price & Availability

XStream Logic has not yet announced products based on the company's Dynamic Multistreaming architecture. For more information, visit the company's Web site at [www.xstreamlogic.com](http://www.xstreamlogic.com).

COPYRIGHT 2000 Cahners Publishing Company

COMPANY NAMES: Xstream Logic Inc.--Product development

GEOGRAPHIC CODES/NAMES: 1USA United States

DESCRIPTORS: Company product planning; Company technology development; Microprocessor

EVENT CODES/NAMES: 331 Product development

PRODUCT/INDUSTRY NAMES: 3674000 (Semiconductor Devices); 3674124 (Microprocessor Chips)

SIC CODES: 3674 Semiconductors and related devices

NAICS CODES: 334413 Semiconductor and Related Device Manufacturing

FILE SEGMENT: CD File 275

02409123      SUPPLIER NUMBER: 62710910      (THIS IS THE FULL TEXT)  
**IBM PAVING THE WAY TO 0.10 MICRON : First Copper, Then SOI, Now Low-k and  
E-Beams. (Company Business and Marketing)**  
Diefendorff, Keith  
Microprocessor Report, 14, 5, 1  
May, 2000  
ISSN: 0899-9341      LANGUAGE: English      RECORD TYPE: Fulltext  
WORD COUNT: 6553      LINE COUNT: 00532

TEXT:

As it has many times in the past, IBM is once again blazing the trail to next-generation IC processing way ahead of the rest of the semiconductor industry. Two years ago (see MPR 9/14/98-msb, "IBM Delivers on Copper Promise With 750-400"), IBM rocked the industry with its leap to copper interconnects--a feat most other vendors are still scrambling to match. A year later, IBM made another startling announcement: it would move its mainstream logic processes to silicon-on-insulator substrates (see MPR 8/24/98-02, "SOI to Rescue Moore's Law"). The company has now made good on that promise by shipping an SOI-based PowerPC processor, code-named IStar, to its AS/400 group. Then, just last month, on April 3, IBM announced yet another giant technological leap, this time to a low-k process (k (less than) 3.0) using a spin-on polymer dielectric, called SiLK by developer Dow Chemical. Copper, SOI, and SiLK will be the baseline materials for IBM's 0.13-micron-generation CMOS-9S process, which will enter production next year.

As if copper, SOI, and low-k weren't sufficient to prove its prowess, on March 2 IBM announced a breakthrough in electron-projection lithography (EPL). This development, which dramatically boosts e-beam-stepper throughput, could potentially render unnecessary the enormously expensive extreme-ultraviolet (EUV) optical steppers that are currently the odds-on favorite for next-generation lithography (NGL). This IBM development could lead to a commercial EPL stepper from partner Nikon by early 2003, opening the door to billion-transistor chips.

While leadership in any one of these technologies would be impressive, IBM's command of all of them is almost unbelievable. Only Motorola, which until last year was a partner of IBM, has so far managed to get copper processors into mass production (see MPR 11/16/98-04, "G4 Is First PowerPC With AltiVec"), but even Motorola is still well behind IBM on copper manufacturing. Other companies have claimed use of "low-k" dielectrics, but these companies are mostly referring to fluorine-doped silicon-dioxide materials with dielectric constants only about 10% lower than conventional SiO<sub>2</sub>. A few companies have also claimed to be working on SOI, but none that we know of (besides IBM) is yet to the stage of seriously considering it for volume mainstream manufacturing. And while a few companies are funding industry consortia research into next-generation lithography, most will simply wait until NGL tools become broadly available from traditional equipment suppliers.

In Conscious Pursuit of a Risky Strategy

IBM could just be blowing smoke, tooting its technology horn more loudly than other semiconductor vendors to gain the appearance of a technology leader. But history does not support this theory. Over the years, IBM has demonstrated a clear pattern: invest heavily in research and development on aggressive new technologies; announce them when they're ready; ram them into volume production; then disseminate the technology to the rest of the industry while moving on to new technologies before the crowd catches up.

Bijan Davari, IBM Fellow, vice president of IBM's Semiconductor Research & Development Center in East Fishkill (NY), and the mastermind of IBM's semiconductor R&D strategy, admits this strategy involves some risks. For one thing, the development of advanced processes is extraordinarily expensive. For another, proprietary processes are not consistent with low-cost manufacturing. On the one hand, IBM would like to maximize the return on its investment by keeping its technology to itself to use as a competitive weapon. On the other hand, it realizes that it cannot afford to be out on a technology limb by itself. IBM needs other semiconductor

manufacturers to adopt its technology so that the equipment industry will invest in developing the reliable low-cost, high-throughput tools that IBM needs for high-volume chip production.

Davari's plan to resolve this dilemma is twofold: stay ahead and partner with other companies. If IBM can stay ahead of the industry, he argues, it opens a window of time during which the company can exploit an advanced technology before others catch up. During this period, Davari says that IBM Microelectronics garners a significant amount of business building for its customers' parts that simply cannot be built by any other vendor.

If IBM stays far enough ahead, then even after this period of exclusivity, its intellectual property will still have enough residual value to be licensed to close partners and, eventually, to the rest of the industry. IBM then plows these licensing revenues back into process development to fund its efforts to stay ahead. Also, IBM allows selected partners, such as UMC and Infineon, to pitch in to help defray development costs in return for earlier access to some of IBM's advanced technologies (see MPR 2/14/00-02, "IBM, Infineon, UMC Gang Up On 0.13").

#### Capacitance, the Microprocessor's Worst Enemy

The transition time of a signal on a wire in an IC is proportional to the product of the wire's resistance ( $R$ ) and its capacitance ( $C$ ). Thus, lowering  $R$  and  $C$  reduces signal delay. Furthermore, the noise that a signal accumulates as it propagates through a wire is related to the degree of capacitive coupling to adjacent signals. Thus, reducing capacitance both reduces signal delay and improves signal integrity.

Unfortunately, capacitance does not scale with process shrinks. The capacitance a signal encounters is proportional to the area of adjacent parallel conductors and inversely proportional to the thickness of the insulator between them. As process dimensions shrink, wires get shorter, reducing  $C$ , but they also get closer together (which increases  $C$ ) and narrower (which increases  $R$ ). Thus, the net effect of process scaling is to leave the  $RC$ -delay component roughly the same, or to make it somewhat worse. So, as process dimensions shrink and transistors speed up,  $RC$  interconnect delay becomes an increasingly large component of overall circuit delay. Furthermore, capacitive coupling of noise among signal lines gets worse, because vertical wire thickness generally isn't reduced by the same scale factor as horizontal line widths and spaces (thickness is usually maintained to keep resistance to a minimum).

$RC$  delay and noise coupling have not always been huge problems. In 0.25-micron and larger processes, transistors largely dominated circuit delays, and wires were far enough apart that only long parallel buses created serious noise problems. But at 0.18 micron, things change: interconnect delays and noise become more significant problems. And at 0.13 micron, unless something is done, these problems become serious obstacles to continued circuit-speed increases.

IBM made a step-function improvement in this situation for its 0.22-micron (CMOS-7S) and 0.18-micron (CMOS-8S) process generations when it introduced copper as the interconnect material. Copper has about 40% lower resistivity than the aluminum alloy used previously, a fact IBM exploits to build thinner interconnect layers--which have less capacitance--without increasing resistance. Figure 1, which compares IBM's 0.18-micron copper CMOS-8S interconnect system with Intel's 0.18-micron aluminum P858 system, clearly illustrates the advantage of copper in this respect.

Although this improvement is substantial, as dimensions shrink further, to 0.13 micron and beyond, and the wires get even closer, capacitance once again becomes a limiting factor. This time, however, no new conductor material will come to the rescue. Silver, the only material more conductive than copper (at normal temperatures), is only slightly more so (about 5%). Fortunately, manufacturers have one more handle on capacitance: the permittivity of the insulator, also called the dielectric constant.

#### Finding the Least-Worst Alternative

The interlayer dielectric (ILD) material used by most manufacturers today is silicon dioxide ( $\text{SiO}_2$ ), which has many ideal physical properties for this purpose. As a glass, it is mechanically solid, allowing it to provide good support for the interconnect layers and to form a tight hermetic seal from the environment. Silicon dioxide is chemically inert and thermally stable, making it compatible with the silicon substrate, with all types of interconnect materials, and with high-temperature manufacturing

steps. In addition, the material offers low leakage currents and high breakdown voltages. It also has excellent adhesion and is amenable to planarization using chemical-mechanical polishing (CMP).

Unfortunately, silicon dioxide doesn't have such ideal electrical properties. Pure SiO<sub>2</sub> has a dielectric constant ( $k$ ) of about 4.0; including overcoats necessary in the manufacturing process, silicon-dioxide insulation typically delivers a  $k$  in the range of 4.3-4.5. Some manufacturers, including IBM in CMOS-8S and Intel in P858 (see MPR 1/25/99-06, "Intel Raises the Ante With P858"), use a fluorine-doped silicon dioxide called FSG (fluorosilicate glass) or SiOE FSG is attractive because it has manufacturing properties similar to pure SiO<sub>2</sub>; unfortunately, it improves the  $k$  by only about 10%. The improvement in a copper-interconnect environment is even less (about 6%), because less fluorine must be used to remain compatible with copper.

While many materials have lower  $k$  than pure or fluorinated SiO<sub>2</sub>, all other known insulating materials are inferior to SiO<sub>2</sub> with respect to their thermal, mechanical, or chemical properties, making them more difficult to use in manufacturing, or less desirable in the final product. It is an intrinsic property of low- $k$  materials, for example, that they also have a low modulus--that is, they are soft. IBM spent several years identifying possible candidates, which are shown in Table 1, and deciding which had the fewest drawbacks--or at least had only problems IBM thought it could tackle.

Another criterion IBM imposed on its search for a low- $k$  material was the requirement that it be extensible. For example, IBM knows that in the future (below 0.10 micron) it will have to adopt porous insulating materials to get a  $k$  closer to 2.0. These advanced porous materials are likely to be "spin-on" materials, as opposed to being applied with a plasma-enhanced chemical-vapor-deposition process (PECVD), as is silicon dioxide. Porous materials, however, will not be ready for manufacturing for several years. Therefore, for this generation, IBM wanted a spin-on material that would be compatible with future tool sets, allowing a smooth transition to porous materials when the time arrives.

#### Plastic Dielectric Is Smooth As SiLK

The dielectric material that IBM finally settled on for its 0.13-micron CMOS-9S process belongs to a class of materials known as aromatic thermosets, specifically an organic polyarylene-ether resin sold commercially by Dow Chemical under the brand name SiLK (see sidebar). Pure SiLK has a dielectric constant of 2.62; including overcoats, SiLK delivers a  $k$  of around 3.0, about 25% better than FSG and more than 30% better than pure silicon dioxide.

Although Dow will sell SiLK to the industry, it will not be easy for other manufacturers to follow in IBM's footsteps. Ron Goldblatt and Jim Ryan, key contributors to IBM's low- $k$  effort, point out that they had to develop a number of new techniques to integrate SiLK into IBM's copper process, which is shown in Figure 2.

One problem with SiLK is that, unlike SiO<sub>2</sub>, it etches at the same rate as resist, a characteristic that makes it incompatible with the traditional copper dual-damascene process flow. To solve this problem, IBM developed a dual hardmask consisting of two dissimilar layers. The dual-damascene pattern is first etched into the hardmask layers, then transferred to the SiLK dielectric. Other techniques had to be developed to compensate for SiLK's low modulus (4% that of SiO<sub>2</sub>) and poor thermal conductivity (15% that of SiO<sub>2</sub>). IBM's techniques involve, among other things, special structures for supporting the interconnect layers and bond pads, changes in design rules to account for SiLK's different etch properties, and optimization of the barrier films to guard against copper contamination.

Solving this latter problem was one of the most challenging for the IBM team. Integrating a new dielectric material into a conventional aluminum metal system isn't an easy task, even though aluminum is chemically benign and its characteristics are thoroughly understood. But integrating a completely new nonoxide-based dielectric with copper--which is highly contaminating and understood much less well--is a far more challenging task. Motorola has previously disclosed progress toward integrating a porous inorganic dielectric ( $k = 2.0$ ) with its copper-metal system (see MPR 5/31/99-msb, "Motorola Takes Capacitance to New Low"), but it admits that much work remains to be done to put that dielectric into

production. IBM is the only company we know of that has cleared all the hurdles of integrating a low-k dielectric into a high-volume-production copper process.

This fact may shed light on IBM's strategy to make an early jump to copper in its 0.22-micron (CMOS-7S) and 0.18-micron (CMOS-8S) processes. The move to copper was criticized by many industry experts, who thought the move was unnecessarily aggressive. Intel, for example, argues that at 0.18 micron, it can achieve equivalent performance just by adding a low-k dielectric (SiOF) to its existing aluminum metal system. While that may be true, IBM now has two generations of copper-manufacturing experience under its belt and thus has a stable next-generation interconnect platform from which it can make the move to a true low-k material. By procrastinating, copper-naysayers will be facing a giant step up when they move to 0.13-micron lithography, copper interconnects, and a new dielectric material all in one generation.

IBM intends to deploy copper and SiLK across its entire process family, including its less-expensive foundry processes. The company has announced that in 3Q00 it will offer a design kit for Cu-11, a 0.13-micron CMOS-8SF ASIC with 40 million wireable gates. It expects to begin sampling the part in 1Q01 and be in full production by 3Q01. In this part, IBM will exploit the low resistance and capacitance offered by its copper/SiLK interconnect system to pack wires more closely together, doubling the number of wireable gates over the previous CMOS-7SF part. IBM says the embedded DRAM array in this part will be 40% denser and 25% faster than the embedded DRAM in its previous CMOS-7SF ASIC.

The embedded-DRAM cell in next-generation CMOS9SF ASICs will be based on yet another IBM innovation: a vertical access transistor that is self-aligned with a buried strap into the trench capacitor. The vertical transistor eliminates the problems associated with continual shrinking of the gate length, thereby allowing a smaller cell size. The technique, which IBM described at the International Electron Devices Meeting (IEDM) last December, reduces the size of a DRAM cell by 25% compared with conventional cells.

#### Fast Interconnects and Fast Transistors

Copper interconnects and low-k dielectrics are all about reducing wire delay. And, to the extent that wire delay is a limiting factor in circuit speed, they do improve the situation significantly. To quantify the gain, IBM performed a complete 3D parametric extraction to simulate signal propagation through four different metal/dielectric systems. As Figure 3 shows, the simulation of a 200-micron M3 wire showed a 37% reduction in wire delay for copper/SiLK over aluminum/SiO<sub>2</sub>--not including any indirect gains from reduced capacitive noise coupling (crosstalk). Because of the conservative assumptions used in the simulation, IBM says it sees even better performance in real silicon than is predicted by the simulation: measurements indicate that copper alone provides up to 20% improvement rather than the 11% predicted by these simulations.

Of course, if wire delay isn't a limiting factor, then the gains predicted in Figure 3 will not result in faster overall circuits. Intel, in its campaign to defend its decision to forgo copper in P858 (see MPR 2/28/00-02, "Processors Penetrate Gigahertz Territory"), says it knows how to rebias the design to be transistor-delay dominated, eliminating potential gains from interconnect speedups. We find this argument unconvincing, however; while this unnatural technique may minimize wire delay, it is not clear that it results in faster circuits. In fact, Texas Instruments found that signal-propagation speed is optimal when gate delay and wire delay are balanced (within a clock cycle), and we estimate that in most 0.18-micron processors today, wire delays and gate delays contribute equally to circuit speed--notwithstanding Intel's techniques.

Moreover, since gate speed increases much more dramatically than interconnect speed when going from one process generation to the next, wire delay will rapidly become the dominant delay term. By the time we reach 0.10 micron, or maybe even 0.13 micron, most of the 37% speed gain IBM predicts from copper and SiLK will manifest itself in higher processor frequencies. The remainder of the problem--gate delays--IBM is attacking aggressively with SOI and lithography.

#### "Industry Must Go to SOI," Says IBM

According to Ghavam Shahidi, manager of IBM's SOI program, scaling of bulk CMOS becomes extremely difficult below 0.13 micron, primarily due to



short-channel effects. As transistor channel length shrinks, parasitic factors, which at long channel lengths were insignificant, become dominant. Loss of gate control (and transistor gain), high gate-overlap capacitance, subthreshold leakage, and tunneling, among other problems, conspire to eliminate the speed gains that have accompanied all previous process shrinks.

Although a few tricks remain at 0.13 micron to counteract some of these problems, at 0.10 micron and below they become unmanageable. Intel, for example, in a paper presented at last December's IEDM, described a notchedpoly technique that undercuts the gate poly to reduce overlap capacitance. IBM says it shies away from such stopgap solutions, however, because they do not scale well to shorter channel lengths. IBM says that even at 0.18 micron, notched poly is more trouble than it's worth. The problem is that ultraprecise control over the etch is required to achieve consistent gate lengths, but such precise control is difficult because of factors such as the proximity of other structures, which create unavoidable local variations in the effectiveness of the etch.

Solutions to other short-channel problems are equally hard to find. At extremely tiny dimensions, manufacturing tolerances simply cannot be kept tight enough to adequately control source/drain doping profiles, for example. And some effects simply cannot be eliminated, even if manufacturing tolerances are perfect. For example, as transistors shrink, the critical charge required to upset SRAM cells and dynamic nodes is lowered. Below 0.13 micron, soft errors induced by charged particles become a big problem, putting a limit on how far these devices can be scaled. But, thanks to the isolation provided by its buried-oxide layer, SOI has a naturally immunity to such disturbances and thus has a much lower soft-error rate (SER) than short-channel bulk processes.

IBM's research into these issues has convinced Shahidi and Davari that there is simply no viable solution to scaling problems in general, save for one: silicon-on-insulator. SOI offers many advantages over bulk CMOS, which we detailed in our 1998 SOI article (see MPR 8/24/98-02, "SOI to Rescue Moore's Law"). The advantage Shahidi cites in defense of IBM's bold assertion that the industry must move to SOI, however, is that SOI offers another knob for controlling the shape of the channel. As Figure 4 shows, the silicon layer above the buried oxide--whose thickness can be precisely controlled--allows source/drain profiles that cannot be created otherwise, solving many of the short-channel problems. This extra knob also allows the creation of unique device structures with characteristics precisely matched to specific circuit needs.

IBM has been building SOI-based microprocessors for some time now, and through that effort it has gained considerable insight into SOI's properties. This experience, according to Davari, has given IBM increasing confidence that SOI is the right strategic path. Simple experiments, such as rendering the same PowerPC design in both bulk CMOS-7S and 7S-SOI, have demonstrated a raw speedup of more than 20% across a 7-sigma variation in channel lengths. Other experiments indicate that redesign to utilize the variable-threshold voltages ( $V_t$ ) and deeply stacked gates made possible by SOI (and impossible in bulk CMOS) can achieve speed gains of 50%, and sometimes more. If these results carry through to volume production, which IBM says they will, just on the basis of SOI alone (independent of copper and low-k), IBM could be one full generation ahead of the industry in process speed while using the same lithography.

IStar, PA-8700 Debut in SOI

Proving that it isn't kidding about its move to SOI, IBM quietly revealed that it is shipping production 540MHz CMOS-7S-SOI processors, code-named IStar, to its AS/400 group. (IStar is a PowerPC-compatible processor with modifications for use in AS/400s.) The company did not say when IStar-based E-Server systems would be available, but historically it takes several months to put server systems into production, indicating availability early in the second half of this year.

IStar, which IBM first described at ISSCC in February of 1998, is essentially the same design as its predecessor, Pulsar, which operates at 450MHz in bulk CMOS-7S. A direct comparison between IStar and Pulsar provides powerful evidence in support of IBM's claim of 25% speed boost due strictly to SOI, without redesign.

In fact, this comparison may underestimate the gain from SOI. Since Pulsar has been in production for some time, its Leff is probably being

pushed more aggressively than that of the new IStar. If true, IBM probably still has enough headroom to push IStar's speed closer to 600MHz, making it 33% faster than Pulsar. Whether the company will make this move depends on how quickly it intends to follow with a CMOS-8S2 version. (CMOS-8S2 is an SOI-only process.) According to IBM's data, shown in Figure 5, 8S2 is 20% faster than 7S-SOI at nominal channel lengths and 33% faster at aggressive Left: Thus, an 8S2 version of IStar should easily coast to 700MHz.

In an announcement that shocked everyone, including IBM, HP disclosed on April 11 the details of an 800MHz PA-8700 processor, which will be available in systems by 1 H0 1. While the 8700 announcement was expected, the disclosure that it would be built in a copper SOI process was a surprise. Although HP didn't officially announce the fab for the 8700, IBM is the only vendor on the planet with a production-worthy copper SOI process. Thus, the mystery of who is building PA-RISC chips these days is now pretty much settled. In fact, the HP-IBM linkage is so transparently obvious that IBM execs are probably more than mildly upset with HP for preempting their official SOI AS/400 announcement.

HP is not the only company looking to IBM for process technology. Sun recently confirmed our suspicions that its MAJC-5200 (see MPR 10/25/99-04, "Sun Makes MAJC With Mirrors") will be built by IBM rather than by its long-time UltraSPARC partner, Texas Instruments. The 5200 is now entering production in 0.22-micron copper CMOS-7S, but it will soon move to 0.18-micron copper CMOS-8S, and eventually to SOI.

Also, at Microprocessor Forum last October, Compaq said that its **21464** would be constructed in a 0.13-micron copper low-k SOI process (see MPR 11/15/99-msb, "Alpha **21464** targets 1.7GHz in 2003"). Furthermore, rumors persist that Compaq is on the verge of announcing a deal with IBM to produce copper Alphas, probably the 21264, probably in CMOS-8S. Given Compaq's Microprocessor Forum statements, we suspect it is also negotiating for access to SOI-based CMOS-8S2 and CMOS-9S for its 21364 and **21464**. Such a deal would be a good move for Compaq and would give us a more favorable outlook on the future of Alpha.

These revelations by Compaq, HP, and Sun represent strong votes of confidence from the industry's top performance leaders for IBM's copper/SOI/low-k process roadmap.

#### Seeking Unlimited Resolution

While IBM pushes hard on the materials front with copper, low-k, and SOI, it is not ignoring the lithography front. Today, for 0.18-micron processes, nearly all manufacturers rely on optical projection lithography using deep ultraviolet (DUV) light at a wavelength of 248nm. But this wavelength is just adequate to image the smallest features on a 0.18-micron chip while maintaining adequate depth of field for high-yield, high-volume production.

To go below 0.18 micron requires a number of resolution-enhancement techniques (RETs), such as off-axis illumination (OAI), strong-phase-shift masks (PSMs), optical proximity correction (OPC), and increased numerical-aperture lenses. Using these techniques, 248nm optical lithography can be pushed to serve the 0.13-micron generation--barely. The 1999 International Technology Roadmap for Semiconductors (ITRS99) calls for a transition to 193nm steppers during the 0.13-micron generation, which, with RETs, will suffice down to 0.10 micron--again, barely. For the 0.10-micron generation, the ITRS99 calls for another wavelength reduction, to 157nm. This time, RETs will allow 157nm steppers to serve down to 0.07 micron, but beyond that DUV isn't workable because, among other factors, lenses just become too opaque.

Therefore, during the 0.07-micron generation, the ITRS99 calls for a transition to a next-generation lithography (NGL) approach. There are four basic candidates for NGL: extreme-ultraviolet lithography (EUVL), X-ray lithography (XRL), electron-projection lithography (EPL), and ion-projection lithography (IPL). Although there is no industry consensus on which is the best approach, the majority of activity and investment over the past few years has been on EUVL, which, at a wavelength of 13.4nm, is suitable for as long as anyone reading this article is likely to care.

Intel has been the primary driving force behind EUVL, and it has formed an industry consortium, called the LLC, to help develop the technology. Three of the major national laboratories--Lawrence Livermore, Sandia, and Lawrence Berkeley--carry out the majority of the work for the LLC, which, surprisingly, includes AMD and Motorola. Sematech also

contributes to the LLC's efforts.

This lithography roadmap, however, is not without problems. Chief among them is cost. Today, a single-column 248nm optical stepper costs \$8 million to \$12 million, and a large fab typically has a couple dozen of them. Replacing this equipment with 193nm steppers will be enormously expensive--not to mention the additional cost of RETs, which is also high. Some industry analysts believe that optical lithography will simply be too expensive for 0.10-micron processing, due both to the cost of equipment and to the poor yields that some expect as a result of narrower and narrower process windows. To turn around and repeat this exercise for 157nm DUV just a couple of years later would be staggering.

At one time it was hoped that EUVL would be ready for the 0.07-micron generation, possibly eliminating the need for the intermediate 157nm DUV step. This, however, does not appear to be feasible. The progress on EUVL has been excruciatingly slow, and the cost of the EUVL systems is likely to be higher than originally projected.

#### E-Beams to the Rescue

Meanwhile, IBM has been plugging away with its EPL research. For many years, the company used e-beam direct write (EBDW) to quickly turn bipolar chips for its mainframes. Initially, its Gaussian-beam EBDW steppers, which raster scan the circuit pattern directly onto the wafer at a rate of one pixel per flash, had lousy throughputs of 0.01 wafers per hour. During the 1980s--when feature sizes were 2 microns and there were fewer than 101(degrees) pixels on a 5mm wafer--IBM coaxed throughput upward to 20 wafers per hour with several-hundred-pixel-per-flash shaped-beam projectors.

The writing speed of these EBDW tools, however, did not keep pace with the Moore's Law rate of pixel growth, and it became clear that throughput would never be adequate for today's high-volume production, which will soon require writing 1013 pixels on a 200mm wafer. (Today's DUV steppers routinely achieve throughputs of 80-100 wafers/hour, and EUVL steppers--which are similar except for their use of mirrors rather than lenses--should have similar throughput.)

But IBM did not give up on e-beams. The company's latest breakthrough is the development of a practical e-beam projection-lithography (EPL) system, which uses mask projection analogous to that used in optical lithography. EPL is attractive as an NGL candidate because its resolution, for all intents and purposes, is unlimited. Both EPL and EUVL are capable of being extended to the 0.035-micron generation and beyond. EPL, however, has never been used with any success in semiconductor manufacturing because of practical limitations, primarily that of limited field size.

One source of problems, says IBM Fellow Hans Pfeiffer, is that electrons are charged particles, and they repel each other (Coulomb interactions). This effect tends to blur the image at high beam intensity. Moreover, while an EPL projected field can be larger than that of a EBDW system, it is still much smaller than most chips, requiring the field to be scanned over a considerable distance to cover the chip. Deflecting the beam very far, however, introduces off-axis aberrations that defeat attempts to contain Coulomb interactions.

To increase throughput in spite of these problems, IBM had to find a way to apply massively parallel pixel projection across a large field without creating distortion. For this feat it developed a novel magnetic lens system that minimizes off-axis aberrations by electronically shifting the optical axis of the lenses in sync with the beam. As Figure 6 shows, this creates a variable curvilinear axis for which the system is named PREVAIL (projection reduction exposure with variable-axis immersion lenses).

IBM is no longer the only company that believes in e-beams. It was apparently able to convince Nikon, the largest supplier of optical steppers today, that its system was a viable NGL contender. Together, the two companies have constructed a proof-of-concept EPL system, shown in Figure 7, that employs a high-emittance, high-numerical-aperture e-beam source along with a silicon stencil mask and a proprietary distortion-correction system. The prototype system, which currently delivers a 12.8(micro)A beam current during each 100(micro)s pulse, has been used to successfully demonstrate 0.08-micron lithography over a 5mm-wide field without significant loss of resolution, as Figure 8 shows.

IBM expects to coax its PREVAIL alpha-tool performance to a 15gA beam

current, delivering 10 million pixels per flash over a 7mm-wide field, which would support a throughput of 35 wafers per hour. On the strength of this prototype system, Nikon says it will build a commercial stepper for deployment in 2003.

Although Pfeiffer admits that EUVL systems will have some advantages over EPL systems, he says that production EPL steppers can be delivered earlier than production EUVL systems with competitive throughput, and that EPL steppers could cost even less than today's DUV optical steppers. If this is true, it would certainly make a compelling case for EPL as the industry's next-generation lithography system. IBM is currently investigating methods for extending EPL to 0.05 and 0.035 micron without sacrificing throughput.

#### Firing on All Cylinders

IBM has always been recognized by the industry as a technology leader. But other semiconductor companies have come to realize that technology is an incredibly important weapon in the microprocessor business--no architectural, microarchitectural, or circuit design innovation is likely to have even close to the impact of a half-generation lead in semiconductor technology. And conversely--nothing is likely to be as devastating as a half-generation technology lag. With such high stakes, other companies have also been investing heavily in advanced semiconductor process development, making us wonder just how long IBM could maintain its preeminent position at the top of the IC-process totem pole.

Despite heavy investment by other companies, however, IBM recently seems to be pulling even further ahead. The string of announcements over the past two years has been truly impressive. While other companies nibble around the edges of next-generation process problems, IBM takes giant bites out of them. Copper, SOI, plastic dielectrics, and e-beam lithography are big bites. But each move the company makes, while unquestionably aggressive, seems to be well justified.

Moreover, they are synergistic. While each technology is valuable in its own right, the combination is awesome. Together, copper, SOI, and SiLK support new design methods capable of producing chips that are easily twice as fast as could be built with a conventional bulk aluminum/SiO<sub>2</sub> 0.13-micron process. Other companies will eventually follow in IBM's footsteps, some willingly, some not. At this point, however, unless other companies are being incredibly secretive, IBM appears to be a good two years ahead of the rest of the industry.

IBM's technology lead is not going unrecognized. Nearly every major semiconductor vendor is actively trying either to license technology from IBM or to emulate it. UMC and Infineon, for example, have just entered into a major technology agreement with IBM (see MPR 2/14/00-02, "IBM, Infineon, UMC Gang Up On 0.13"). Motorola and AMD have joined forces to develop copper HIP6 and future processes that are likely to include SOI and low-k dielectrics (see MPR 8/3/98-msb, "Motorola, AMD Swap Technology"). We expect that even Intel, although it is forced to go slow because of its enormous volumes, will eventually follow IBM's lead, as it has done previously on such IBM innovations as shallow-trench isolation.

Moreover, nearly all high-performance processor design houses (except Intel, Motorola, and AMD) are beating down IBM's door to gain access to its advanced processes. Plans by Intel-partner HP for the PA-8700, TI-partner Sun for MAJC, Samsung/API-partner Compaq for Alpha, and startup Transmeta for Crusoe (see MPR 2/14/00-01, "Transmeta Breaks x86 Low-Power Barrier") are all strong endorsements of IBM's semiconductor technology. Every company in the world that is building a performance- or power-critical microprocessor or SOC knows instinctively that IBM is the place to look for the best technology. They also know, however, that it is the most expensive place to look. IBM is proud of its technology and is not ashamed to ask a premium price for it.

One next-generation technology on which IBM has been notably silent is the issue of 300mm (12-inch) wafers. John Kelly, the general manager of IBM Microelectronics, has stated that IBM doesn't intend to be the first company to use the foot-wide wafers. It's not a surprise, however, that IBM would be slow to adopt 300mm wafers. Although 300mm wafers are important from a fab-capacity point of view (300mm wafers carry 2.5 times as many chips as 200mm wafers), they do not directly contribute to performance, power, logic density, or reliability, which are IBM's primary concerns. Besides, the company does not intend to lag far behind the industry on

300mm. Davari says IBM will begin the transition to 300mm wafers during the 0.13-micron CMOS-9S and CMOS-8SF generations, putting it only slightly behind leaders Intel (see MPR 6/21/99-msb, "Intel Commits to 300-mm Wafers") and UMC (see MPR 1/24/00-04, "Hitachi, UMC Jump on 12' Wafers").

IBM's strategy to stay ahead of the rest of the industry on technology is a bold one, if not an extremely risky one. To stay on this fast-moving treadmill, IBM cannot afford to stumble. A single misstep, such as falling into a losing-technology rat hole, could easily throw IBM off the treadmill, which runs too fast to get back on. To guard against such risks, IBM is attempting to follow a very well thought out long-range roadmap and to distribute the risks by working in parallel on multiple technology fronts. So far, the strategy is working, but it will require extreme vigilance to continue this strategy ad infinitum. IBM is silent on when or from where its next process advancement will come. But given its strategy and its past performance, it is a safe bet that East Fishkill researchers have something up their sleeve.

Materials	k	Process
Silicon Dioxide	3.9-4.5	PECVD
Fluorosilicate Glass (FSG)	3.2-4.0	PECVD
Polyimides	3.1-3.4	Spin-on
HSSQ	2.9-3.2	Spin-on
Diamond-Like Carbon	2.7-3.4	PECVD
Carbon-Doped SiO <sub>2</sub>	2.7-3.3	PECVD
Parylene-N	2.7	CVD
Benzocyclobutenes	2.6-2.7	Spin-on
Fluorinated Polyimides	2.5-2.9	Spin-on
MSSQ	2.6-2.8	Spin-on
Aromatic Thermosets	2.6-2.8	Spin-on
Fluorinated DLC	2.4-2.8	PECVD
Parylene-F	2.4-2.5	CVD
Teflon AF	2.1	Spin-on

#### RELATED ARTICLE: A Really Low k

IBM has announced that it will use SiLK resin from Dow Chemical in its next-generation 0.13-micron CMOS-9S process, which will enter volume production next year.

SiLK is spin-on aromatic hydrocarbon polymer with a dielectric constant of 2.62. SiLK is stable at temperatures of up to 450(degrees)C, allowing it to withstand the rigors of the semiconductor manufacturing process. The new material has an etch selectivity of 20:1 and can be etched with standard O<sub>2</sub> /N<sub>2</sub> plasma. It is compatible with either aluminum or CVD- or electroplated-copper metal systems. With a toughness of only 0.62MPa-m<sup>1/2</sup>, however, SiLK is softer and less adhesive than traditional silicon-dioxide interlayer dielectrics (ILDs), making it difficult to planarize with conventional chemical-mechanical polishing (CMP)--a problem IBM had to work around.

Dow and IBM are now working together on ultra-low-k (k ? 2.0) porous dielectrics for 0.10 micron and beyond as part of the National Institute of Standards and Technology's advanced technology program.

For more information on SiLK go to [www.silk.dow.com](http://www.silk.dow.com).

SiLK Property	Value
Dielectric Constant (k )	2.62
Leakage Current	3.3x10 <sup>-10</sup> A/cm at 1mV/cm
Breakdown Field	4mV/cm
Glass Transition Temp (Tg)	(greater than)490(degrees)C
Thermal Stability	(greater than)450(degrees)C
Young's Modulus	2.7 GPa
Toughness	0.62 MPa-m <sup>1/2</sup>
Film Stress	45 MPa
Moisture Uptake	0.25% @ 80% RH, 25(degrees)C
Thermal Conductivity	0.18 W/mK
Crack Growth Rate in Water	(less than)10 <sup>-11</sup> m/sec

COPYRIGHT 2000 Cahners Publishing Company

COMPANY NAMES: International Business Machines Corp.--Research

GEOGRAPHIC CODES/NAMES: 1USA United States

DESCRIPTORS: Research and development; Company technology development

EVENT CODES/NAMES: 310 Science & research  
PRODUCT/INDUSTRY NAMES: 3674000 (Semiconductor Devices)  
SIC CODES: 3674 Semiconductors and related devices  
NAICS CODES: 334413 Semiconductor and Related Device Manufacturing  
TICKER SYMBOLS: IBM  
FILE SEGMENT: CD File 275

2/9/14 (Item 8 from file: 275)

DIALOG(R) File 275:Gale Group Computer DB(TM)  
(c) 2004 The Gale Group. All rts. reserv.

02358406 SUPPLIER NUMBER: 58248504 (THIS IS THE FULL TEXT)  
**Compaq Chooses SMT for Alpha : Simultaneous Multithreading Exploits  
Instruction- and Thread-Level Parallelism. (Compaq Alpha 21464 ) (Product  
Information)**

Diefendorff, Keith  
Microprocessor Report, 13, 16, NA  
Dec 6, 1999

ISSN: 0899-9341 LANGUAGE: English RECORD TYPE: Fulltext  
WORD COUNT: 5376 LINE COUNT: 00433

TEXT:

As it climbs rapidly past the 100-million-transistor-per-chip mark, the micro-processor industry is struggling with the question of how to get proportionally more performance out of these new transistors. Speaking at the recent Microprocessor Forum, Joel Emer, a Principal Member of the Technical Staff in Compaq's Alpha Development Group, described his company's approach: simultaneous multithreading, or SMT. Emer's interest in SMT was inspired by the work of Dean Tullsen, who described the technique in 1995 while at the University of Washington. Since that time, Emer has been studying SMT along with other researchers at Washington. Once convinced of its value, he began evangelizing SMT within Compaq. His efforts have apparently paid off, as Compaq has officially adopted SMT for the Alpha 21464 (see MPR 11/15/99, p. 13), code-named EV8, which is due to appear in systems in 2003. That Compaq is talking about this processor three full years in advance indicates great confidence in SMT technology as well as a strong desire to establish that Alpha has a future.

SMT processors are similar to conventional super-scalar out-of-order processors, but they have additional hardware resources that allow them to interleave the execution of multiple instruction streams, or threads, onto the execution units, as Figure 1 shows. By more fully utilizing the execution units in this way, SMT processors achieve higher sustained throughput and improved tolerance of memory latency.

The Debate: ILP or TLP

Even with 100 million of them on a chip, transistors are not free-yet. Hence, the question persists of how to deploy them in a way that maximizes performance. One alternative is to use them just to build larger on-chip memories, as Intel has done with the new Pentium III (see MPR 10/25/99, p. 1). This approach is effective, but only up to a point, beyond which little is gained from adding more cache. At that point, performance becomes limited by the speed of the processor core.

Given a full complement of on-chip memory, increasing the clock frequency will increase the performance of the core. One way to increase frequency is to deepen the pipeline. But with pipelines already reaching upwards of 12-14 stages, mounting inefficiencies may close this avenue, limiting future frequency improvements to those that can be attained from semiconductor-circuit speedup. Unfortunately this speedup, roughly 20% per year, is well below that required to attain the historical 60% per year performance increase. To prevent bursting this bubble, the only real alternative left is to exploit more and more parallelism.

Indeed, the pursuit of parallelism occupies the energy of many processor architects today. There are basically two theories: one is that instruction-level parallelism (ILP) is abundant and remains a viable resource waiting to be tapped; the other is that ILP is already tapped out, and it's time to move on to the richer vein of thread-level parallelism (TLP). TLP proponents point to the rather depressing history of ILP progress. Over the past 10 years, processors have grown from simple

single-issue machines with fewer than 1 million core transistors to four-wide out-of-order behemoths with 10-million-transistor cores. At the same time, however, sustained ILP has done little better than double. Professor John Hennessy of Stanford in his keynote speech at the Forum showed data indicating that while the theoretical ILP of an assortment of SPEC95 benchmarks ranges from about 18 to 150 IPC, practical four-wide out-of-order super-scalar processors rarely achieve even 2 IPC. ILP pessimists further assert that progress will be more dismal in the future, as diminishing returns will more severely curtail ILP gains.

#### No Lack of Ideas to Use Transistors

ILP proponents counter that with just a few more tens of millions of transistors, mixed with a little compiler magic, they can unleash this wealth of ILP. According to this group, radical models of execution that could not be considered in past are becoming feasible. HP and Intel with Itanium (see MPR 10/6/99, p. 1) are depending on static instruction scheduling by a compiler, predication, and very large register files to achieve a step-function increase in ILP. Hal's Sparc64 V (see MPR 11/15/99, p. 1) is using trace processing and super-speculation to achieve high ILP. Optimists like Yale Patt at the University of Texas and John Shen at Carnegie-Mellon believe that such advanced super-scalar techniques will allow ILP to scale with transistor count, ultimately enabling 16- or 32-wide processors with sustained ILP of 10 IPC or more on general-purpose applications. Even if they succeed in extracting this degree of parallelism, such processors will all have one thing in common: physically large monolithic cores. Many will also have staggeringly complex control mechanisms. Although transistors will be plentiful enough to implement such machines, physics will surely intervene to enforce its immutable rule that large things are slow things. Furthermore, design and verification will become ever more difficult and time consuming. These realities, along with less confidence in ILP, have motivated IBM with Power4 (see MPR 10/6/99, p. 11) and Sun with MAJC (see MPR 10/25/99, p. 18) to shift their attention from ILP to explicit thread-level parallelism. These companies are using their transistors to build chip multi-processors (CMPs). They believe it is wiser to keep processor cores small and fast, by limiting their issue widths, while relying on the parallelism between independent program threads to achieve higher performance.

A major drawback of both high-ILP processors and CMPs is that they suffer from poor transistor utilization when the workload doesn't match the processor. High-ILP processors speculate poorly or leave function units idle when faced with programs having inherently low ILP. Similarly, CMPs must leave entire processors idle when enough threads aren't available.

#### Enter SMT

Simultaneous multithreaded processors are a cross between wide-issue super-scalar processors and fine-grain-multithread processors (see MPR 7/14/97, p. 13). Fine-grain multithreading (FMT) was first implemented by Seymour Cray in the peripheral-processing unit of the CDC 6600 (circa 1964), then again in the late 1970s in Denelcor's HEP, and more recently in Tera's MTA. FMTs maintain state information for several active threads, and on each cycle they issue one instruction from a different thread. The advantage of this technique is that it fills pipeline bubbles created by dependencies on long latency operations (e.g., memory accesses) with instructions from known-independent threads. This is far easier and more effective than trying to fill bubbles by ferreting out and reordering independent operations from a single thread. If FMT were straightforwardly extended to super-scalar issue, as Figure 2 shows, it would address the problem of low temporal utilization of execution units (pipeline bubbles), but the problem of low spatial utilization (empty execution slots) would remain, due to intrathread dependencies. Simultaneous multithreading, however, allows instructions to be selected for issue from any ready thread, as Figure 1 shows. In this way, SMT processors can fill unused execution slots with useful work.

The real beauty of SMT is that as threads execute, the machine can dynamically reallocate execution resources on the basis of the mix of parallelism in the workload. A single thread with a high degree of ILP can utilize the full resources of the machine for maximum speed; alternatively, resources can be distributed among several threads to achieve high throughput, even in the face of low ILP. Indeed, any combination of workload types can execute concurrently, with performance limited only by

the total available resources.

SMT's ability to exploit parallelism in a wide variety of workloads produces consistently high execution-unit utilization, a fact that enables designers to consider wider super-scalar designs than could be justified on ILP alone. Although Emer counts himself in the camp of ILP-optimists and says that EV8 would have been eight-wide even without SMT, he is less sanguine about ILP beyond that which is exploitable by an eight-wide processor. With SMT to take up the excess, however, even wider machines might be effective.

#### Not That Different From Wide Super-scalar

Conceptually, SMTs are similar to wide-issue dynamically scheduled processors, as Figure 3 shows. In fact, no new control mechanisms are needed to issue instructions from multiple threads. The traditional register-renaming scheme, for example, avoids false dependencies (register name conflicts) between threads in the same way it does within a single thread: by mapping architectural registers from the active thread onto the processor's pool of physical registers. This is not to say that no additional hardware is required for SMT. Thread identifiers, for example, must be appended to each instruction so thread-specific operations, such as branch prediction and virtual address translation, can be performed as instructions flow through the pipeline. Also, some processor resources must be duplicated so that state information (registers, program counter, etc.) can be maintained separately for each active thread context.

Other hardware, such as that required for recovering from branch mispredictions, handling program exceptions, maintaining precise interrupts, returning from subroutines, and retiring instructions in order, must either be replicated for each thread or shared, which requires more complex bookkeeping logic.

While the aforementioned additions are required to achieve proper function, even more hardware is probably needed to carry the heavier load of multiple threads. Instruction queues must be deeper, and more registers must be available in the renaming pool. Caches, translation-lookaside buffers (TLBs), and branch-history tables (BHTs), should also be larger, be more associative, and have more ports. And because the SMT's execution units are shared among several simultaneous threads, their number and symmetry may have to be increased to prevent contention.

While these additional hardware resources do not themselves add much complexity beyond that found in a conventional super-scalar processor, they do add size. To prevent this size increase from impacting cycle time, steps must be taken that do indeed increase complexity. The caches, for example, may have to be partitioned into multiple smaller banks; the register files and execution units may also have to be partitioned, as they are in the 21264 (see MPR 10/28/96, p. 11); and the pipeline may have to be lengthened, putting pressure on the branch predictor, rename registers, and reorder buffer.

Even though SMTs require incremental hardware to support each thread, an SMT capable of running four simultaneous threads, for example, would be nowhere near four times larger than a single-thread super-scalar of the same issue width. Two things account for this economy: first, SMT threads exploit hardware that would otherwise be sitting idle; second, the statistical variations in multiple threads running asynchronously prevent excessive contention for some hardware. Thus, a good deal of hardware—the execution units and caches, for example—can be effectively shared, avoiding hardware increases for each thread. Indeed, Compaq says that EV8's resources are sized for a single-thread and that additional SMT threads are treated as opportunistic. Future processors, however, may indeed have beefed-up resources to reduce conflicts.

#### Instruction Fetch Limits Throughput

Because the SMT has more independent instructions at its disposal (from separate threads), it can issue instructions at a far higher rate than a single-thread processor. This higher issue rate puts severe pressure on the instruction fetcher. In fact, instruction fetch is potentially the most severe bottle-neck in an SMT processor. Therefore, it is necessary to minimize branch mispredictions, to minimize fetching speculative instructions when nonspeculative ones are available, and to have an intelligent mechanism for selecting threads from which to fetch.

Emer described one possible scheme in a paper he co-authored for the Sept/Oct '97 issue of IEEE Micro. In the eight-wide SMT hypothesized in



that paper, on every cycle the instruction fetch unit fetched eight instructions from each of two threads that were not currently processing an instruction-cache miss. Instructions were selected for dispatch from the first thread until either a branch or an end-of-cache line was encountered, at which time instructions were selected from the second thread.

The two threads were selected using an Icount feedback technique. This technique prioritizes fetch from threads that currently have the fewest instructions in the front end of the pipeline. The theory behind Icount is that it gives the highest fetch priority to the fastest-moving threads and maximizes interthread parallelism by maintaining an even distribution of instructions from different threads in the instruction queues. Icount also prevents thread starvation, since threads with the fewest instructions in the pipeline are the first to get new fetch cycles.

Icount scheduling has the fortuitous characteristic of very low hardware cost; all that's required is a simple up/ down counter for each thread and some comparators to select the two threads with the smallest counts. In the Micro paper, the researchers found Icount to be more effective than alternative schemes that sought to fetch from threads in ways that minimized branch mispredictions or load delays.

But Does It Work?

Apparently so. Although no one has yet built an SMT, simulations show it to be promising. On the hypothetical eight-wide machine in Emer's Micro paper—which had six ALUs (four of which can load or store) and four FPUs—Emer reported a speedup of slightly more than 2x for four threads over one. The speedup held for both multi-programmed single-thread applications and for single multithreaded programs. To simulate the worst case for multiprogramming (most potential interthread contention), the same application was executed for all four threads. Applications were selected from the SPEC95 and Splash2 benchmark suites.

Performance gains flattened abruptly for more than four threads; eight threads showed no appreciable benefit over four. Presumably, four threads were able to saturate the execution resources of the hypothetical machine, limiting further gains. This result was probably influential in Compaq's decision to limit the eight-wide EV8 to four active threads.

To support more than four simultaneous threads, EV8's fetch, dispatch, and issue widths would probably have to be increased along with the number of execution units. Since EV8 is slated for a 0.125-micron process and a 250-million-transistor budget, we doubt it was concern over transistor count that limited the width. Instead, it was probably the complexity and cycle-time implications of going beyond eight-wide super-scalar. It could also have been the enormous demand SMT puts on memory: just to support four threads, EV8 will have a direct multichannel interface to RDRAM main memory, and, although Compaq has not stated this, it will probably have more than 3M of on-chip L2 cache.

At the Forum, Emer presented additional simulated-benchmark results, further illustrating the speedup achievable by an SMT processor. As Figure 4 shows, with a multi-programming workload of mixed integer and floating-point benchmarks, four-way SMT had nearly 125% higher throughput on four threads than on one. On multithreaded programs, four-way SMT achieved better throughput by an average of 75%, as Figure 5 shows. The SPECfp95 benchmarks in this suite were automatically decomposed into threads, and Emer says a manual decomposition may produce better results. These results are impressive, considering the modest amount of hardware required to support three additional threads. Compaq claims that for EV8 the additional silicon area for its four-thread SMT core above the base eight-wide super-scalar core is less than 10%. In comparison, doubling the silicon area of a single-thread processor typically boosts performance by less than 50%—and that percentage is trending down. We know of no other EPIC or advanced super-scalar approach that could double the performance of an eight-wide super-scalar Alpha processor for less than double the silicon. Thus, the approximately 2x speedup Emer reported would seem to make SMT a real bargain.

It is important to remember, however, that Emer's benchmarks measure speedup when thread-level parallelism is present. In real systems, however, sometimes TLP will not be present. Thus, in practice, the speedup from SMT will, on average, be less than Emer's benchmark results show.

It is also important to note that there is a big difference in

complexity and area between a four-wide and an eight-wide super-scalar. Therefore, these results cannot be used to compare EV8 with the alternative of, for example, two four-wide super-scalar processors on a CMP (chip multiprocessor). In his Micro paper, however, Emer reported that a CMP with two cores, each having roughly half the resources of the hypothetical eight-wide SMT, showed similar speedups for two threads, but it fell well short of SMT's four-thread performance. The SMT is also likely to have better single-thread performance than the CMP when ILP is present.

#### The Pesky Matter of Software

As is frequently the case with techniques to speed up processors, SMT is not without software issues. Although SMT executes single-thread programs with no difficulty, problems creep in when you try to use its multithreading capability. For multiprogramming workloads (workloads comprising multiple individual programs running simultaneously), the problems are tractable; the software implications are minor and restricted to the operating system. For this case, the OS simply needs to prioritize threads and to keep the most important thread contexts resident on the processor. Multiprogramming speedup, however, is important today only in server environments that are currently served by symmetric multiprocessors (SMPs).

To fully justify SMT, however, it is necessary to also take advantage of single-program multithreading. To enable this, programs must be decomposed into multiple independent threads that the SMT can execute in parallel. This requires two things: the presence of thread-level parallelism in the program and the ability to find and expose it.

Unfortunately, techniques for automatically decomposing programs into parallel threads are in their infancy. Guri Sohi at the University of Wisconsin is pursuing multiscalar techniques in which a single thread is decomposed into mini-"tasks" according to the program flow graph; multiple task sequencers then use aggressive control and data-value speculation to execute these tasks in parallel. Former graduate student Scott Breach has shown that enhanced SMT hardware can be used to run these mini-tasks in parallel.

But how effective compilers will be in automatically creating parallel threads from a single program remains to be seen. Today, the burden of parallelizing programs remains a largely manual process. To make matters worse, debugging multithreaded programs is notoriously difficult—a fact that deters many programmers. Although multithreading is becoming a more accepted style of programming, especially with Java, today most programs are still single threaded, and most programmers are still poorly trained to code for explicit parallelism. This obstacle could prevent SMTs from realizing their full potential for several years. Perhaps by 2003, when EV8 systems are due to appear, things will have changed.

The architectural abstraction that Compaq has adopted for programming EV8 is that of a CPU with four thread-processing units (TPUs), as Figure 6 shows. This abstraction creates a programming model of SMT as a sort of virtual CMP. In fact, the SMT is functionally similar to CMP in many ways. For example, both share data among threads without going off chip, both exploit thread-level parallelism, and both can switch thread contexts in about the same amount of time.

One difference between the two, one that Compaq's abstraction makes clear, is that SMT threads share data at the L1 without the overhead of the cache-coherency actions required by CMPs with separate L1s. This feature gives SMTs the potential for slightly finer-grain threading and tighter coupling between threads. On the other hand, because they share data at the L2—and do not share L1s, BHTs, TLBs, execution units, or anything else—CMPs provide a higher degree of thread isolation; that is, the performance of one thread is less dependent on the characteristics of other threads than it would be on an SMT. This isolation may be an advantage in some situations, such as in critical real-time applications.

To make the TPU model work, one problem Compaq had to eliminate was the problem of spin loops. Whenever multiple threads cooperate, mechanisms are needed to synchronize threads, communicate between threads, lock shared resources, and protect critical software sections. These functions are normally accomplished in software by low-level semaphore operations that involve putting the processor into spin loops while polling for semaphore changes. Spin loops in an SMT, however, are a disaster because they consume one of the TPUs while performing no real work.

To circumvent this problem, Compaq devised a method for putting a thread to sleep and waking it when a given memory location changes. Instructions are not fetched or issued from a sleeping thread, allowing other active threads to utilize more of the processor's resources. The scheme was inexpensive to implement, as it relies on the existing load-with-lock/store-conditional semaphore mechanisms already in the Alpha architecture and the cache-coherency mechanisms that already exist to detect cache-line modifications.

Won't Affect Cycle Time, Right?

According to Emer, SMT need not lengthen cycle time. Emer believes that the cycle time of a CPU should be set according to the highest speed that the ALU can evaluate and forward results to subsequent instructions. The pipeline length should then be established by dividing execution into stages no longer than the ALU cycle time. But SMTs need more registers and thus longer operand-read times than a super-scalar. To prevent these factors from impacting the cycle time, it is very likely that at least one additional pipeline stage will be required, which would add to the branch-mispredict penalty. Other SMT-specific resources, such as more instruction-completion writeback ports, could impose additional stages.

As a result, it is likely that a single-thread application will not perform as well on an eight-wide SMT as it would on a super-scalar of similar design. This loss of single-thread performance, if indeed it is only one pipeline stage, probably amounts to only a few percent. If SMT turns out to have other resources or control complexities that add more pipeline stages or increase cycle times, the net benefit of SMT will be less clear. But Emer sees no reason to expect any cycle-time penalties or any more than one or two extra pipeline stages.

Another potential performance limitation is resource contention among threads. Even in a 0.125-micron process, execution units will not be completely symmetric, and not every structural hazard will be eliminated. Even worse, contention for the caches, the BTB, and the TLB could increase miss rates or, in the worst case, cause severe thrashing. Assuming these resources are sufficiently associative, thrashing should be avoidable, but cache miss rates will definitely go up, due to increased conflicts. SMT's greater ability to tolerate memory latency should compensate to some degree-but to what extent remains to be seen. Compaq says it has seen cases of positive interference, such as prefetching system code, but these cases are probably the exception rather than the rule.

Alternatives Abound

With transistor budgets soon to exceed 100 million transistors per chip, a host of architects with ideas on how to spend those transistors has emerged. The most popular ideas being espoused for general-purpose microprocessors, aside from SMT, include advanced super-scalar processors (e.g., trace processing, superspeculation, and multi-scalar), EPIC (explicitly parallel architectures), and CMP (chip multiprocessors).

These ideas are not necessarily mutually exclusive and could conceivably be used in combination. In the near term, however, sheer size and complexity will preclude most combinations. Longer term, with say a billion-transistor budget, nearly any combination could, in theory, be built. But many hybrids will not bear fruit, regardless of the transistor budget. SMT is likely to be incompatible with some of the advanced super-scalar techniques. These techniques, for example, frequently depend on speculation. But SMT and speculation both vie for the same resources, and both stress the fetch unit to achieve their goal.

SMT is probably even less compatible with EPIC than it is with advanced super-scalars. Although Intel has alluded to the possibility of eventually adding multithreading to future IA-64 implementations, it is not clear that move will be feasible. SMT depends, by its very nature, on the dynamic-scheduling hardware that is present in super-scalars but is completely lacking in EPIC. Adding these mechanisms on top of EPIC would risk massive complexity, and it would defeat one of its central tenets. Furthermore, the result may be disappointing. EPIC, using predicated execution, attempts to fill idle function units with speculative operations from the current program thread. To the extent it succeeds in this objective, EPIC would reduce the effectiveness of SMT by usurping the very execution units on which SMT thrives, as Figure 7 shows.

Since the techniques are not always synergistic, SMTs will likely end up facing advanced super-scalar and EPIC processors in the market. Against

these techniques, SMT will have the powerful advantage of being able to evoke either thread-level or instruction-level parallelism at will. This advantage will materialize only when enough total parallelism is available, but this flexibility will allow the SMT to perform well in many situations where these other techniques would fail miserably.

SMT will have the disadvantage, however, that in single-thread environments, programs must be explicitly written to expose thread-parallelism. If programs do not migrate to multithreaded construction, then SMT's additional resources will go for naught, and its single-thread performance is likely to be inferior to one of the other techniques using equivalent resources.

In server environments, which are usually heavily multiprogrammed, this disadvantage will not come into play. But even in a multiprogrammed or multithreaded environment, SMT will be of little benefit if individual programs have high ILP. In such a case, the execution units will be kept busy by the high-ILP thread, leaving few execution slots for other threads.

SMT, CMP Square Off SMT's most serious long-term challenge will probably come from CMPs, which have some compelling advantages of their own. A CMP core, because it is typically smaller and less ILP-aggressive than an SMT core, is likely to achieve a higher frequency and/or have a shorter, more efficient pipeline. If ILP turns out to be limited or to be hard to exploit with wide-issue machines-and there is precious little hard evidence to the contrary-then CMPs, which can also play the thread-level-parallelism card, might perform as well as an SMT.

If performance is similar, then CMP construction wins. Building one small, simple core and replicating it along with a shared L2 is a far simpler and more expeditious task than designing a large, complex, monolithic core. In addition, CMPs introduce the potential for using partially-good die. This possibility can reduce manufacturing scrap, thereby reducing the average manufacturing cost of a CMP die.

Because an SMT shares more resources among threads, it will probably have a physically smaller die than an equivalent performance CMP. But this advantage may be less than it seems. For one thing, given upcoming transistor budgets, sharing resources may not save enough silicon to be worth the control complexity needed to do so. Second, the high execution-unit utilization of SMTs could create longer queue delays and longer latencies that would require additional hardware in the SMT to ameliorate. Third, utilization is naturally higher and therefore less of a problem on narrow-issue CMP cores. SMTs, in a sense, create an artificially low utilization situation by starting out with an excessively wide-issue engine. In the future, CMP and SMT techniques might create an interesting marriage. If low utilization is a problem even for modest four-wide super-scalar CMP cores-which, with a throughput of less than 2 IPC, would seem to be the case- then a simple four-wide/two-thread SMT core might eliminate the problem. Arraying this core in CMP fashion might provide a simple path for scaling beyond the four active threads that are the limit of an EV8-class eight-wide SMT. Architects of IBM's Power4 CMP (see MPR 10/6/99, p. 11) have already expressed a possible interest in multithreading for the future.

Putting multiple EPIC or advanced-super-scalar processors on a chip will be another way to exploit ILP and TLP; the question is whether there is enough ILP to justify using these more complex cores. Although this option may not be realistic in the near term-say over the next three to four years, while transistor budgets are limited to a measly 100 million to 250 million transistors per chip-in the long term it could pose a powerful alternative to SMT.

In the meantime, the one incontrovertible advantage of SMT-and the characteristic that makes it attractive over all other known forms of advanced super-scalar, EPIC, CMP, or combinations thereof-is its unique ability to shift resources on the fly between ILP and TLP at a very fine grain. The ultimate value of this advantage, however, will depend heavily on software evolution.

To go beyond servers, either something like multimedia must drive up the use of multiprogramming in PC environments, or a much broader range of applications must move to multithreaded construction. This move could happen quickly if compiler techniques evolve to automatically create parallel threads, or if Java-which already has multithreaded API classes and background tasks-takes hold. If either event happens over the next

three years, we may see more vendors adopting the clever technique of SMT.

For multiprogrammed server environments, however, SMT is readily applicable. And Compaq says the programs used in many of Alpha's key application areas, such as data warehousing, graphics rendering, and government super-computing, are already multithreaded. Assuming that Compaq remains committed to Alpha, and doesn't let annoying details such as IC process and system design stand in its way, SMT should provide a solid basis for the company to retain Alpha's long-standing performance title over all comers.

For More Information

"Simultaneous Multithreading: Maximizing On-Chip Parallelism," Tullsen, Eggers, and Levy, ISCA95.

"Exploiting Choice: Instruction Fetch and Issue on an Implementable Simultaneous Multithreaded Processor," Tullsen, Eggers, Emer, Levy, Lo, and Stamm, ISCA96.

"Converting Thread-Level Parallelism to Instruction-Level Parallelism via Simultaneous Multithreading," Lo, Eggers, Emer, Levy, Stamm, and Tullsen, ACM Transactions on Computer Systems, August 1997.

"Simultaneous Multithreading: A Platform for Next-Generation Processors," Eggers, Emer, Levy, Lo, Stamm, and Tullsen, IEEE Micro, October, 1997.

COPYRIGHT 1999 Cahners Publishing Company

COMPANY NAMES: Compaq Computer Corp.--Products

GEOGRAPHIC CODES/NAMES: 1USA United States

DESCRIPTORS: Product description/specification; Microprocessor; Processor architecture

EVENT CODES/NAMES: 330 Product information

PRODUCT/INDUSTRY NAMES: 3674124 (Microprocessor Chips)

NAICS CODES: 334413 Semiconductor and Related Device Manufacturing

TRADE NAMES: Compaq Alpha **21464** (Microprocessor)--Design and construction

FILE SEGMENT: CD File 275

2/9/15 (Item 9 from file: 275)

DIALOG(R)File 275:Gale Group Computer DB(TM)

(c) 2004 The Gale Group. All rts. reserv.

02349930 SUPPLIER NUMBER: 57588372 (THIS IS THE FULL TEXT)

Alpha **21464** Targets 1.7 GHz in 2003. (Compaq details plans for Alpha EV8 processor) (Company Business and Marketing)

Microprocessor Report, 13, 15, NA

Nov 15, 1999

ISSN: 0899-9341 LANGUAGE: English RECORD TYPE: Fulltext

WORD COUNT: 489 LINE COUNT: 00040

TEXT:

Determined to maintain leadership performance well into the next century, Compaq disclosed plans for its futuristic **21464** processor at last month's Microprocessor Forum.

Work on the forthcoming Alpha design, codenamed EV8, is already under way, but the chip is not scheduled to appear in systems until early 2003.

According to Compaq's Joel Emer, the **21464** will achieve single-thread performance leadership using an eight-way superscalar processor core running at speeds of up to 1.7 GHz. The new core's instruction-reordering capabilities will be enhanced significantly over those of the current **21264** to accommodate the greater issue width. As a result, Emer expects perclock performance to nearly double compared with the **21264**.

The high clock speed will be delivered by a 0.13-micron CMOS process with advanced features such as copper, low-k dielectrics, and SOI. Compaq did not name the fab, but we expect the **21464**, like current Alpha chips, will be built by Samsung and possibly another foundry. Emer said the design will achieve clock speeds of 2.0 GHz and up, but these speeds will require more advanced IC process technology.

To further boost performance, the **21464** will implement four virtual processors on the chip, using a technique called simultaneous

multithreading (SMT). This method allows instructions from up to four separate threads to share a common CPU core, filling dead cycles in one thread with unrelated instructions from another thread. Emer said SMT will increase performance by as much as 2x in a multi-threaded environment, as is common in servers.

The system interface of the 21464 will be similar to that of the 21364 (see MPR 10/26/98, p. 12), which is due to appear in systems in early 2001. Like that chip, the 21464 will have a large on-chip L2 cache, several Rambus channels for main memory, and four additional ports for accessing other processors' memory. Presumably, the cache size and memory bandwidth will be increased from the 21364's to accommodate the more powerful core, but Compaq declined to provide additional details.

In 2002, the 21464 will compete against Intel's Madison, a 0.13-micron version of McKinley. We expect Madison to achieve similar clock speeds, and its IA-64 design may offer a performance advantage on single-threaded programs.

Both chips are likely to deliver in excess of 130 SPECint95 (base). The 21464, however, could have an edge in servers, due to its multithreaded design; we expect the McKinley/Madison core will be single-threaded.

When discussing processors so far in the future, the biggest question is whether the vendors will be able to deliver on schedule. We won't know that answer for quite some time, but, for now, Compaq's announcement shows it is not backing down from IA-64 in the performance race.

COPYRIGHT 1999 Cahners Publishing Company

COMPANY NAMES: Compaq Computer Corp.--Product development

GEOGRAPHIC CODES/NAMES: 1USA United States

DESCRIPTORS: Microprocessor; Company product planning

EVENT CODES/NAMES: 331 Product development

PRODUCT/INDUSTRY NAMES: 3674124 (Microprocessor Chips)

NAICS CODES: 334413 Semiconductor and Related Device Manufacturing

TRADE NAMES: Compaq Alpha EV-8 (Microprocessor)--Product development

FILE SEGMENT: CD File 275

?t s2/9/16

2/9/16 (Item 10 from file: 275)

DIALOG(R)File 275:Gale Group Computer DB(TM)

(c) 2004 The Gale Group. All rts. reserv.

02338696 SUPPLIER NUMBER: 56025118 (THIS IS THE FULL TEXT)

**Power4 Focuses on Memory Bandwidth. (new IBM architecture) (Product Development)**

Diefendorff, Keith

Microprocessor Report, 13, 13, NA

Oct 6, 1999

ISSN: 0899-9341

LANGUAGE: English

RECORD TYPE: Fulltext

WORD COUNT: 6484

LINE COUNT: 00501

TEXT:

IBM Confronts IA-64, Says ISA Not Important

Not content to wrap sheet metal around Intel microprocessors for its future server business, IBM is developing a processor it hopes will fend off the IA-64 juggernaut. Speaking at this week's Microprocessor Forum, chief architect Jim Kahle described IBM's monster 170-million-transistor Power4 chip, which boasts two 64-bit 1-GHz five-issue superscalar cores, a triple-level cache hierarchy, a 10-GByte/s main-memory interface, and a 45-GByte/s multiprocessor interface, as Figure 1 shows. Kahle said that IBM will see first silicon on Power4 in 1Q00, and systems will begin shipping in 2H01.

No Holds Barred

On this project, Big Blue is sparing no expense. The company has brought together its most talented engineers, its most advanced process (0.18-micron copper silicon-on-insulator), and its best packaging, reliability, and system-design know-how. The sheer scale of the project indicates that IBM is mindful of the threat posed by IA-64 (see MPR 5/31/99, p. 1) and signals that the company is prepared to fight for the server market that it considers its birthright.

After years of building their own processors, IBM, HP, and others have been forced to watch as systems based on commodity Intel microprocessors have chipped away at their market. HP recognized the futility of continued resistance and threw in the towel. But IBM sees that with more and more of the critical system-performance features moving onto the processor, the loss of control over the processor silicon would rob it of the ability to assert its superior technology and to differentiate itself from the pack.

Although the IBM PC Company has already elected to go with IA-64 for its Netfinity servers, IBM apparently believes it cannot strategically afford to do the same for its high-end (high-margin) server businesses, where it makes a large portion of its revenues today and which it expects will grow rapidly along with the Internet. Therefore, the company has decided to make a last-gasp effort to retain control of its high-end server silicon by throwing its considerable financial and technical weight behind Power4.

After investing this much effort in Power4, if IBM fails to deliver a server processor with compelling advantages over the best IA-64 processors, it will be left with little alternative but to capitulate. If Power4 fails, it will also be a clear indication to Sun, Compaq, and others that are bucking IA-64, that the days of proprietary CPUs are numbered. But IBM intends to resist mightily, and, based on what the company has disclosed about Power4 so far, it may just succeed.

#### Looking for Parallelism in All the Right Places

With Power4, IBM is targeting the high-reliability servers that will power future e-businesses. The company has said that Power4 was designed and optimized primarily for servers but that it will be more than adequate for workstation duty as well. The market IBM apparently seeks starts just above small PC-based servers and runs all the way up through the high-end high-availability enterprise servers that run massive commercial and technical workloads for corporations and governments.

Much to IBM's chagrin, Intel and HP have also aimed IA-64 at servers and workstations. IA-64 system vendors such as HP and SGI have their sights set as high up the server scale as IBM does, so there is clearly a large overlap between the markets all these companies covet. Given this, it is surprising that they have come to such completely different technical solutions.

Intel and HP have concluded there is still much performance to be found in instruction-level parallelism (ILP). Hence, they have mounted an enormous effort to define a new parallel instruction-set architecture (ISA) to exploit it (see MPR 5/31/99, p. 1). Evidently, they expect a significant speedup from machines that can issue six or more instructions per cycle (any less wouldn't justify a new ISA).

IBM, in contrast, believes the place to find parallelism in server code is not at the instruction level but at the thread level and above. It doesn't believe there's enough ILP in individual threads of server code to fill a large number of instruction-issue slots. Even if there were, IBM says that EPIC-style architectures like IA-64 are contraindicated. Although high-ILP processors may reduce processor busy time, IBM points out that they do nothing to reduce processor wait time, which is the far larger problem. In fact, it says EPIC architectures exacerbate this problem by burdening the memory system with a large number of conditionally executed instructions that are eventually discarded.

#### Dynamic Scheduling Is Better, Says IBM

Power4 engineers cite a number of arguments in favor of dynamic scheduling over EPIC-style static scheduling for servers. One issue is cache misses; dynamic machines constantly remake the instruction schedule, thereby avoiding many pipeline stalls on cache misses. EPIC machines, because of their in-order execution and static instruction groupings, are less adaptive. EPIC does allow the compiler more freedom to boost loads, and a register scoreboard like the one in Merced allows some run-time adjustments, but cache misses can be hard to predict at compile time and EPIC machines will generally take less advantage of run-time information than reordering superscalar machines.

Another issue IBM raises is the impracticality of code profiling. According to IBM, profiling large server applications is often difficult, and the results not that valuable. But EPIC compilers rely heavily on profiling information to schedule predication and speculation. Wen-Mei Hwu,

speaking at last year's Microprocessor Forum, spelled out several other EPIC-compiler challenges. IBM believes many of these will not be solved for a long time.

If EPIC compilers for traditional code are a challenge, dynamic just-in-time compilers (JITs) for Java will be a nightmare. EPIC compilers must search a large code window to discover ILP and must perform complex code transformations to exploit predication and speculation. Thus, EPIC compile time can be long, making it hard to amortize at run time. Java performance is a serious issue for IBM, which is committed to Java for server applications and has the second-largest cadre of Java programmers in the world, next to Sun. Sun probably agrees with IBM's concerns about EPIC, as its new MAJC architecture (see MPR 9/13/99, p. 12) has many features that are radically different from IA-64 for just these reasons.

IBM is also concerned that EPIC binaries are too tightly coupled to the machine organization. Although Intel and HP have taken steps to ensure that IA-64 code will function across generations, IBM says that an EPIC instruction schedule is so dependent on the machine organization that, in practice, it will restrict hardware evolution.

But IBM's primary objection to EPIC isn't that it's bad, it's just that it's so unnecessary. IBM sees no difficulty in building dynamically scheduled processors that can exploit most of the ILP in the vast majority of server applications. It also sees no difficulty now or in the future in building dynamically scheduled POWER processors that can fully tax any practical memory system. Therefore, IBM concludes that the memory system is the real determinant of server performance, not the instruction set. Thus, staying with POWER imposes no real penalty and avoids a pointless ISA transition.

#### Chip-to-Chip Interconnect Shares L2

As a result, IBM has focused on system design rather than on instruction-set design. The technology, and most of the silicon, in a Power4 chip is dedicated to delivering data to a large number of processors as quickly as possible. The key element IBM uses to accomplish the task is the shared L2 cache. Power4's on-chip L2 is shared directly by the two on-chip processors and by processors on other chips via a high-speed chip-to-chip interconnect network, as Figure 2 shows.

Details on the physical structure of the network have not yet been disclosed, pending patent applications. Kahle did, however, describe some of its features. The network logically appears to each processor as a simple low-latency bus, while the actual physical network provides the high bandwidth and nearly contention-free throughput of a full crossbar switch, but without the complexity.

The chip-to-chip data paths shown in Figure 2 each include multiple 16-byte-wide point-to-point buses arranged in a ring-like topology that IBM describes only as a distributed switch. The switch is implemented entirely on the Power4 die, with no external chips required.

Physically, each chip-to-chip bus is unidirectional and operates on a synchronous latch-to-latch protocol. The low-voltage signals transfer data at a rate of over 500 MHz, giving each Power4 chip an aggregate sustainable chip-to-chip bandwidth of over 35 GBytes/s. Such high bandwidth keeps the network utilization low, which, according to queuing theory, minimizes network latency. The bus architecture is designed so that when four Power4 chips are located in close proximity and each die rotated 90(THORN), the buses between chips route directly. This keeps the wires very short and therefore allows the buses to be very wide and very fast.

As Figure 3 shows, the shared-L2 cache is divided into three multiported, independently accessible slices. A 100-GByte/s switch connects the L2 slices to the on-chip processors as well as to off-chip processors through the chip-to-chip interconnect ports. A shared-intervention protocol is used to enforce cache coherence and to move data into the L2 on the chip that used it last. The goal of the design is to get the right data into the right L2 at the right time and, from a coherency perspective, make sure it is safe to use.

IBM has not disclosed the size of the L2 cache on each Power4 chip, but, based on 170 million transistors and the floor plan in Figure 3, we estimate that the L2 is about 1.5M. We also expect it to be at least eight-way set-associative, as IBM rarely builds on-die cache of less. Due to the large size of the L2 and the reliability requirements for high-availability servers, the L2 is protected from manufacturing defects by row



and column redundancy and protected from run-time soft errors by ECC.

#### A Memory Bandwidth Behemoth

Each Power4 chip provides an L3-cache port separate from the chip-to-chip ports. The L3 port is 16 bytes wide in each direction and operates at a 3:1 clock ratio, providing over 10 GBytes/s of memory bandwidth. The L3 cache tags are kept on the processor die so cache coherency actions can take place at on-chip cache speeds. From the size of the L3 directory shown in Figure 3, we estimate that each Power4 chip can support up to 32M of external L3 cache.

IBM did not describe the L3 architecture, but Figure 2 shows it to be an inline design. This application is a perfect fit for IBM's embedded-DRAM process, which the company has used before to construct integrated-cache chips. With its latest 0.18-micron CMOS-7SF merged-logic/DRAM process, IBM could easily construct a very large set-associative ECC cache with a high-speed interface to the Power4 chip and an interleaved ECC memory controller to drive the main-memory DRAMs.

To help convert Power4's copious memory bandwidth into low-latency memory accesses, the chip implements eight software-activated prefetch streams. These prefetch streams use spare bandwidth to continuously move data through the memory hierarchy and into the L1. Up to 20 cache lines can be kept in flight at a time. Once the prefetch pipe is filled, the memory system can theoretically deliver new data from main memory to the core every cycle.

#### Chip Multiprocessing Boosts SMP Performance

Placing its bet behind the theory that the most important parallelism in server workloads is above the instruction level, IBM has optimized the Power4 system for shared-memory symmetric-multiprocessing (SMP) performance, as opposed to uniprocessor performance. Instead of spending its transistors on a single monolithic CPU, IBM has opted for two smaller CPUs on each Power4 chip.

The theory is this: above some point, say four instructions per cycle, ILP becomes hard to find, leading to diminishing returns on transistors spent to recover it. This implies that a single monolithic CPU will not scale linearly with transistor count. On the other hand, with efficient data sharing, two processors can be made to scale almost linearly, at least when there are enough independent threads available to keep both cores busy, which is usually the case with server workloads. Thus, for a given transistor budget, two smaller CPUs should outperform one big one.

The key is efficient data sharing, which is what Power4 is all about. The latency and bandwidth between on-chip CPUs and a shared multiported L2 cache can be many times what is achievable with discrete CPUs. For discrete CPUs with separate on-chip L2 caches, shared data must be shuffled between chips across external wires. For discrete CPUs with an external shared L2, every L2 access from both CPUs goes off chip.

In either case, to match the speed of on-chip data sharing, the discrete CPUs would require external buses that are far wider and faster than physics allows. For any given number of wires connecting processors, higher levels of SMP can be achieved with two cores on a chip than with one core. Furthermore, containing all the memory traffic between two CPUs and their L2 on a chip takes an enormous load off the external buses, simplifying the chip-to-chip interconnect.

If this theory is valid, it alone would be enough to justify the chip multiprocessing (CMP) approach IBM has taken with Power4. But CMP has secondary benefits as well. For one, a small simple CPU will generally run at higher clock rates than a large complex one. For another, it is easier to design and replicate a simple CPU than it is to design a complex one.

#### "Simple CPU" Is a Relative Term

For the CMP approach to work, each CPU must be powerful enough to exploit most of the ILP that exists in single threads. Although IBM is not ready to release details of the Power4 CPU microarchitecture, it has given a few clues to suggest that each of Power4's two CPUs will exceed the power of any single microprocessor that exists today.

From the floor plan shown in Figure 3 and the transistor count, we estimate that each CPU core (including L1 caches) contains about 30 million transistors, three times as many as in Pentium III. In addition, each Power4 CPU will run at "over 1 GHz," which probably means at least 1.1 GHz. To achieve these frequencies, IBM set a design goal of 8 to 10 gate delays

between pipeline stages, which, for a RISC- style ISA, probably indicates an integer pipeline of about 10 stages and a load pipeline of about 12; IBM has not confirmed these estimates. We expect each Power4 CPU to be like Power3 and have two fully pipelined double-precision floating-point multiply-add units and two complete load/store units.

Even though IBM disdains IA-64's EPIC approach, it appears to be stealing a page from Intel's playbook. In the same way that Intel usurped RISC principles to implement its x86 CISC architecture in P6, IBM plans to expropriate VLIW principles to implement its RISC architecture in Power4.

IBM only vaguely described the mechanism, but apparently in the early stages of the pipeline, the Power4 CPU groups instructions into VLIW- like bundles. These bundles are dispatched to issue queues, where individual instructions are held until their dependencies are resolved and then issued to the execution units. The pipeline beyond the issue stage is noninterlocked; so, once issued, nothing stops an instruction from completing, but all instructions in a bundle must complete before the bundle is retired.

Unlike conventional superscalar implementations that track individual instructions from dispatch through completion, the Power4 CPU tracks bundles only. According to IBM, this mechanism, along with data-flow sequencing through the noninterlocked pipelines, dramatically simplified the Power4 implementation, cutting the percentage of control logic in half compared with that of the four-issue Power3 design (see MPR 11/17/97, p. 23). This brought the control complexity of Power4 more in line with that of a VLIW machine while preserving the advantages of dynamic scheduling.

IBM said that the out-of-order-completion resources in the Power4 CPU are deep enough to hide the full latency of an L2 cache hit, which is probably 8-10 cycles. Also, to a greater extent than on any previous Power or PowerPC processor, Power4 will exploit the architecturally specified weak-storage-ordering model to reorder memory transactions and hide memory latency.

#### Layering for Frequency

Each Power4 CPU implements the same ISA as IBM's current RS/6000 and AS/400 systems and is also fully PowerPC compatible. IBM did, however, make some improvements that will be invisible to programs. The company is finally acknowledging that some of the complex instructions retained from the original 1990 POWER definition may not have been such great ideas. These instructions hinder the ability to run dynamically scheduled wide-issue processors at high frequency.

Convinced, however, that instruction-set stability is critical to its customer base, IBM didn't take the radical step of expunging these instructions from the ISA. Instead, it has introduced instruction-set layering into Power4. In this strategy, the hardware is optimized for the simple instructions, making no frequency compromises for complex ones. Slightly complex instructions, such as the base-register-update form of loads and stores, are cracked into two simple instructions by the instruction decoders. Moderately complex instructions, such as the string ops, are executed by a simple non-branching microcode engine. The most complex instructions, such as the old POWER instructions that were removed in PowerPC, trap to software emulation routines. In this way, existing binaries run unmodified, but new binaries created by compilers aware of the layering may run faster by exploiting the faster alternatives.

#### Systems of All Sizes

The dual-CPU Power4 chip will serve as the basic building block of a wide range of RS/6000 and AS/400 server systems. The first systems will probably be eight-way SMPs built with four Power4 chips mounted on a multichip module (MCM), as Figure 4 shows. This design point is the sweet spot for Power4 chips, as it utilizes most of the chips' features in their most optimal configuration and balance.

The MCM, designed by IBM for Power4 systems, is not your garden-variety MCM. Since, according to our calculations, each 1.5-V Power4 chip will dissipate over 125 W, the MCM has to dissipate over half a kilowatt. It must also deliver 350 A of noise-free current and transmit thousands of 500-MHz signals among Power4 chips and out to memory.

The solution is a multilayer glass-ceramic substrate with copper interconnect layers. Glass ceramic provides a dielectric constant ( $k$ ) of about 5, 45% lower than conventional alumina-ceramic ( $\text{Al}_2\text{O}_3$ ) substrates ( $k$  (superscript two) 9). The copper interconnect layers offer significantly

lower resistance than the refractory-metal layers (tungsten or molybdenum) used in alumina-ceramic packages.

The processor die are flip-chip mounted into the MCM with a staggering 5,500 100-(micro sign)m C4 solder balls spaced on 200-(micro sign)m centers. Of the 5,500 connections, approximately 2,200 are signal I/Os; the rest provide power and ground. An advanced direct-attach technique improves heat transfer from the silicon to the MCM-package substrate.

As Figure 4 shows, the MCM is mounted on a massive metal carrier that physically attaches it to the motherboard and to its air-cooled heat sink. Since the land-grid-array style package is too large and too expensive to be reflow soldered, we suspect IBM may be using the metallized-particle interconnects (MPI) offered commercially by Thomas & Betts or the CIN::APSE fuzz-button connectors offered by Cinch.

These types of connectors can require as much as 60 grams of force per pad to make reliable electrical contact across such a large package. Thus, with 5,200 pads, the MCM would require a total of about 700 pounds of force to insert. This may explain the thickness of the metal carrier, which must be extremely flat and rigid to evenly distribute that much force while maintaining the necessary planarity. (MPI connectors have a compliance of about 250 microns.)

#### Elastic I/O Connects MCMs

Each Power4 chip has two 16-byte-wide L3/memory buses as well as multiple expansion buses that are routed off the MCM through approximately 3,400 signal pads. The expansion buses, among other things, allow multiple MCMs to be connected together to form larger systems.

IBM calls its expansion buses elastic I/O, due to their unique ability to decouple latency from bandwidth. With traditional buses, the maximum bandwidth of the channel is determined by its latency, which is limited by the end-to-end channel delay and by the worst-case timing skew across the width of the channel. But IBM's elastic I/O uses a low-voltage source-synchronous wave-pipelining technique with per-bit de-skew to eliminate the dependence on channel latency. With IBM's scheme, multiple bits are kept in flight on each wire at the same time, and the per-bit de-skew allows arbitrarily wide buses to operate at high clock frequencies.

The two eight-byte-wide intermodule buses operate at more than 500 MHz, giving each chip a bandwidth of about 8 GBytes/s for a total of about 32 GBytes/s between modules. This bandwidth is probably sufficient to build a four-MCM SMP (32-processor) system with memory-access times sufficiently uniform to support classical SMP workloads without retuning the software for nonuniform memory access (NUMA). In addition to the intermodule buses, the expansion buses include separate buses for I/O and NUMA, bringing the bandwidth of each chip's expansion buses above 10 GBytes/s.

Primarily due to shared-memory bandwidth constraints, neither Power4's nor any other known technology will allow SMP systems to scale beyond a few dozen processors. For applications, such as transaction processing, that are amenable to software partitioning, larger Power4 systems can be constructed in NUMA configurations. Power4 chips have integrated support for large NUMA configurations as well as for IBM's logical partitioning (LPAR) feature, now also supported by Sun in its Enterprise 10000 systems. IBM envisions large Power4 NUMA nodes combined into even larger systems, using the clustering technology developed for its S/390 mainframes and its RS/6000SP multiprocessor systems.

Going the other direction in system size, IBM says it plans to offer the Power4 chip in a single-chip module for small dual-processor SMP servers. Presumably, it could also offer a single-processor system using partially good die. Partially good die is one more advantage of CMP construction. The redundancy of two identical CPUs can, in theory, be exploited to reduce manufacturing scrap, thereby reducing average manufacturing cost. This effect can be substantial for a large die, especially in a new, immature process. But IBM has given no indication it intends to exploit this capability.

#### All Hands to Battle Stations

"Power4" is actually somewhat of a misnomer. The name denotes a part that is simply the next-generation processor in the Power, Power2, Power3 series. But the name vastly understates the size and importance of this project to IBM. Previous Power chips were designed in relative isolation by the small RS/6000 group in Austin. Although viable products, these chips

ran far below industry norms for clock frequencies, and the systems offered no compelling technical advantages. As a result, RS/6000 systems have slipped in market share against Sun, HP, and the myriad Xeon-based systems, disappearing almost completely from the workstation market.

Power4 is an entirely different beast, overpowering all previous Power projects. The only similarity between Power4 and its predecessors is the instruction set. The level of investment is of an entirely different order of magnitude. For Power4, the very best people and technology have been marshaled from every corner of the massive company.

High-frequency circuit-design methods were contributed by IBM Yorktown, which developed the techniques used to design the 637-MHz Alliance G6 mainframe microprocessor, until recently the highest-speed microprocessor shipping from any company. IBM Burlington developed the wave-pipelining technology for the expansion buses. The packaging technology was developed by experts with roots in IBM's Hudson Valley mainframe group. The RS/6000 group in Austin, working jointly with the AS/400 group in Rochester, did the system design. The CPU core was developed by chip architects from the Power3 and Somers group in Austin, with help from IBM's Austin Research Labs and its T.J. Watson Research Labs in Yorktown.

#### Reliable All the Way Down to the Silicon

The CMOS-8S2SOI process was developed in IBM's East Fishkill process-development labs. This 1.5-V seven-layer-metal process is a variation of IBM's 0.18-micron copper CMOS-8S (see MPR 9/14/98, p. 1), which IBM will put into production later this year. The 8S2 derivative has 15% shorter channel lengths ( $L_g < 0.12$  (micro sign)m) and is built on a silicon-on-insulator (SOI) wafer (see MPR 8/24/98, p. 8). According to IBM, the low parasitic capacitance of SOI transistors boosts logic speed by over 25% compared with an equivalent bulk process, while also reducing power consumption.

A major constraint placed upon the development of CMOS-8S2SOI was very high reliability. Most processor manufacturers design their gate dielectrics to a Grade 3 failure-rate specification of 1,000 FITs (failures per billion hours). IBM, however, says this isn't good enough for duty in continuous-availability servers, because internal error-detection features extensive enough to compensate for IC-process-reliability problems would add cost and sacrifice considerable speed. As a result, IBM specifies its processes to a 10-FIT failure rate, two full orders of magnitude better than most companies.

To meet this stringent specification, the 8S2 gate oxide had to be made 3.6 nm thick ( $T_{ox}$  at 1.5 V), 20% thicker than the gate oxide in Intel's 0.18-micron 1.5-V P858 process (see MPR 1/25/99, p. 22), which it will use for Merced and McKinley. IBM had to develop other means to compensate for the losses of transistor drive current and of switching speed that result from the thicker gate oxide. SOI and copper were key to achieving these goals. Copper also improved the reliability of the on-chip interconnects; because the metal is nearly impervious to electromigration, it can sustain higher currents for longer periods without failing.

Even with this level of processes reliability, IBM still included a number of RAS (reliability, availability, and serviceability) features in Power4. IBM isn't ready to reveal all of Power4's RAS features, but it did confirm that the part has traditional features such as ECC on the L2, L3, and main memory. It also said that the Power4 has an independent on-chip full-speed test processor and logic analyzer that can be used during manufacturing and system operation to verify functionality and isolate failures. External testers are simply not viable for gigahertz chips with the amount of on-board logic, memory, and I/O that Power4 has.

#### Systems Still a Long Way Off

Although Power4 looks good at this point, a lot can happen between now and system shipments. Even though IBM feels it has invested enough in Power4 to ensure its success, the company is not invulnerable to technical glitches. IBM has, however, taken a number of risk-management steps, including the fabrication of a large test chip to validate Power4's critical technologies. IBM reported on that chip at this summer's Hot Chips. The company has also scheduled more than ample time between first silicon, due 1Q00, and system shipments, scheduled for 2H01. As a result, technical risk probably isn't IBM's biggest concern.

Cost is also not an issue. In CMOS-8S2, 170 million transistors, half

of them cache, should fit on a 400-mm<sup>2</sup> die. While large, such a die is manufacturable for IBM; it is actually 15% smaller than HP's current PA-8500 (475 mm<sup>2</sup>), which has the same amount of cache. Even assuming \$400 for the MCM and conservative estimates of defect density and wafer costs, the MDR Cost Model projects a manufacturing cost of under \$2,500, hardly unreasonable for an eight-processor module. Besides, in large servers the leverage of the CPU is so enormous that price is rarely an issue.

The real issue for IBM is competition. Compared with today's server microprocessors, of course, there is no contest. Even next year's Foster, Merced, UltraSparc-4, and 21364 aren't likely to be a match for Power4. The real challenge will come from the next generations of these processors, which are due out in late 2001 or 2002. Unfortunately, not enough is publicly known about them to make solid comparisons.

Today, Sun is the most direct competitor for IBM's server business. In the past, Sun has thrived, despite relatively low performance processors, by concentrating on high memory bandwidth and robust multiprocessor systems. With Power4, however, IBM may have Sun outgunned, as it is difficult to imagine anyone creating a system with much higher bandwidth than Power4. If Sun can deliver its 1.5-GHz UltraSparc-5 in late 2001, as planned, it might compete with Power4, but there is some question about Texas Instrument's (Sun's UltraSparc foundry) desire to match IBM's leading-edge IC processes, given its own focus on low-cost DSPs.

Performance-wise, Compaq's Alpha processors are everyone's most feared competitor. The current 667-MHz four-issue out-of-order 21264 is the industry's performance leader. By the time Power4 arrives, the 21264 will have been replaced by the 21364 (see MPR 10/26/98, p. 12). This part will use the 21264 core but boost frequency to 1 GHz with a 0.18- micron process, add a 1.5M on-chip L2, a 6-GByte/s memory port, and 13 GBytes/s of chip-to-chip bandwidth.

In some ways, the system architecture of the 21364 is similar to Power4's. Both employ out-of-order superscalar microarchitecture, large on-chip caches, a dedicated memory port, and a high-speed point-to-point interconnect network between chips. The 21364, however, doesn't offer chip-level multiprocessing, and the topology of the interconnect network is different. The 21364's flat mesh has an elegant symmetry, but it doesn't match Power4's raw bandwidth numbers. Since the topologies are different, however, the bandwidth numbers are difficult to compare.

The **21464**, due out sometime in 2002, will be a multithreaded version of a new core, designed to exploit the thread-level parallelism (TLP) that Power4 exploits with on-chip multiprocessing. CMP and multithreading each have advantages and disadvantages, and it will be interesting to see which approach offers better performance. This assumes, of course, that Compaq will remain committed to Alpha after Merced and McKinley ship, and that it can find a fab capable of matching IBM's.

#### Battle With IA-64 Takes Shape

The most serious competition will surely come from IA-64, not just in HP systems but also from the collective mass of other server vendors that have lined up behind that architecture. The first IA-64 processor, Merced (see cover story), will ship in systems starting in 2H00 and will still be the prevailing IA-64 processor when Power4 arrives in 2H01. Merced is a single six-issue sub-gigahertz processor with a small on-chip L2 and less than a tenth of Power4's chip-to-chip bandwidth, so it isn't likely to match that chip's server performance.

Power4's first real IA-64 challenge will come from McKinley, due in late 2001. Intel and HP say that McKinley will be far superior to Merced. According to some sources, McKinley will run at 1.2 GHz and deliver twice the performance and three times the bandwidth of Merced. McKinley may outrun Power4 on single-thread benchmarks, but it lacks CMP and presumably has far less system bandwidth.

The great unsolved mystery is why Intel/HP and IBM arrived at such polar-opposite solutions. Intel and HP have obviously focused their efforts on exploiting single-thread ILP, with less concern for TLP or memory bandwidth. At the opposite extreme, IBM has focused on massive memory bandwidth and TLP but paid only moderate attention to ILP.

Intel obviously believes there is enough latent ILP lying around to justify a departure from the most dominant architectural franchise in the history of mankind. Intel says it has made the switch to a new ISA at this time to give it a solid platform to which it can later add TLP and

high-bandwidth interfaces. It believes that others will eventually be forced to make this same ISA

transition to avoid leaving a wealth of parallelism on the table.

IBM, on the other hand, clings to a far less pervasive ISA, seeing little rationale for more than minor tweaks. IBM says that memory bandwidth is the limiting factor today and predicts that it will only get worse over time. The company believes that the parallelism achievable with superscalar, multithreading, and multiprocessing can saturate any practical memory system, now and until quantum dots replace transistors. Thus, the whole issue of the ISA is simply a moot point.

Something is obviously amiss; both camps cannot be right. There are a number of possible explanations for the disagreement. One is that the companies are pursuing different markets. This explains some of the differences, but not all. If Intel were solely focused on low-end to midrange industry-standard servers, where price/performance is more important, that would explain the traditional busing and packaging technologies of Merced, and probably McKinley as well.

But this is not a completely satisfactory explanation. Although IBM may be biased more toward the high end than Intel is, HP's target market is right in line with IBM's. Intel and IBM both speak about servers with similar numbers of processors, both talk about high-availability systems, and both are interested in workstations. Given these similarities, it is hard to see how the workloads of the systems Intel and IBM both seem to covet could possibly be large enough to justify such disparate views on computer architecture.

Intel, of course, could have its eye on an even more distant market: PCs. While Intel is initially deploying IA-64 at the high end, where it is easier to flesh out, it may really be optimizing the architecture for future duty in PCs. This explanation makes some sense. After all, IBM may be correct: in servers, memory bandwidth and TLP may matter more than ILP or ISA. But Intel could also be correct: ILP and ISA may be important—just to a different market.

If this explanation is correct, it presents IBM with both a big opportunity and a big problem. With Intel's real attention elsewhere, IBM has a chance to bring its considerable resources and technology to bear exclusively on the server market, possibly establishing a strong market position before IA-64 gains a full head of steam. The risk IBM takes, however, is that the momentum Intel will gain in the broader markets could eventually undermine and overwhelm Power4-based servers, despite any technical superiority.

Another partial explanation for their differences may be Java. IBM is making large investments in Java technology—everything from Java class libraries for server applications to faster compilers and virtual machines. Most Java code is heavily multithreaded, playing directly to the strengths of Power4. Not coincidentally, Sun's Java architecture MAJC (see MPR 8/23/99, p. 13) is also optimized for TLP over ILP. Like Power4, MAJC uses CMP and, like IBM, Sun does not envision high-ILP cores; MAJC is optimized for four-instruction issue.

Power4 Not the End of Line

Even if Power4 is wildly successful in IBM servers, its overall impact on the market will be limited. IBM has no current plans to sell Power4 chips commercially, so other server vendors do not have it as an option. Even if IBM were to sell Power4 chips, it would be too late to derail IA-64. IA-64 appears destined to become the basis of industry-standard servers, and Power4 will always be vulnerable to it.

To prevent encroachment from IA-64, IBM must not only acquire the performance lead with Power4, it must hold it. And this performance lead must be convincing to make its market position unassailable. Of course, IBM is planning for just that. Its roadmap shows frequency increases of 25% every year, with performance growing at three times that rate before jumping dramatically with the mid-decade introduction of a new Power5 design. Considering the strength of the Power4 design and the technology muscle IBM is putting behind it, it may be a long time, if ever, before IA-64 infiltrates the large servers that are at IBM's heart.

COPYRIGHT 1999 Cahners Publishing Company

COMPANY NAMES: IBM Personal Computer Co.--Product development  
GEOGRAPHIC CODES/NAMES: 1USA United States

DESCRIPTORS: Hardware product development; Microprocessor  
EVENT CODES/NAMES: 331 Product development  
PRODUCT/INDUSTRY NAMES: 3674124 (Microprocessor Chips)  
NAICS CODES: 334413 Semiconductor and Related Device Manufacturing  
TRADE NAMES: IBM Power4 (Microprocessor)--Product development  
FILE SEGMENT: CD File 275